



# man pages section 3: Networking Library Functions

---

Sun Microsystems, Inc.  
4150 Network Circle  
Santa Clara, CA 95054  
U.S.A.

Part No: 816-0214-10  
May 2002

Copyright 2002 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright 2002 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



020313@3332



# Contents

---

**Preface 11**

**Networking Library Functions 17**

accept(3SOCKET) 18  
accept(3XNET) 20  
ber\_decode(3LDAP) 22  
ber\_encode(3LDAP) 27  
bind(3SOCKET) 31  
bind(3XNET) 33  
byteorder(3SOCKET) 35  
cldap\_close(3LDAP) 36  
cldap\_open(3LDAP) 37  
cldap\_search\_s(3LDAP) 38  
cldap\_setretryinfo(3LDAP) 40  
connect(3SOCKET) 41  
connect(3XNET) 43  
dial(3NSL) 46  
doconfig(3NSL) 48  
endhostent(3XNET) 50  
endnetent(3XNET) 52  
endprotoent(3XNET) 54  
endservent(3XNET) 56  
ethers(3SOCKET) 58  
fn\_attr\_bind(3XFN) 60  
fn\_attr\_create\_subcontext(3XFN) 61  
fn\_attr\_ext\_search(3XFN) 62

fn\_attr\_get(3XFN) 69  
 fn\_attr\_get\_ids(3XFN) 70  
 fn\_attr\_get\_values(3XFN) 71  
 FN\_attribute\_t(3XFN) 73  
 fn\_attr\_modify(3XFN) 75  
 FN\_attrmodlist\_t(3XFN) 77  
 fn\_attr\_multi\_get(3XFN) 80  
 fn\_attr\_multi\_modify(3XFN) 84  
 fn\_attr\_search(3XFN) 86  
 FN\_attrset\_t(3XFN) 91  
 FN\_attrvalue\_t(3XFN) 93  
 FN\_composite\_name\_t(3XFN) 94  
 FN\_compound\_name\_t(3XFN) 99  
 fn\_ctx\_bind(3XFN) 104  
 fn\_ctx\_create\_subcontext(3XFN) 106  
 fn\_ctx\_destroy\_subcontext(3XFN) 107  
 fn\_ctx\_equivalent\_name(3XFN) 108  
 fn\_ctx\_get\_ref(3XFN) 110  
 fn\_ctx\_get\_syntax\_attrs(3XFN) 111  
 fn\_ctx\_handle\_destroy(3XFN) 113  
 fn\_ctx\_handle\_from\_initial(3XFN) 114  
 fn\_ctx\_handle\_from\_ref(3XFN) 116  
 fn\_ctx\_list\_bindings(3XFN) 118  
 fn\_ctx\_list\_names(3XFN) 119  
 fn\_ctx\_lookup(3XFN) 122  
 fn\_ctx\_lookup\_link(3XFN) 123  
 fn\_ctx\_rename(3XFN) 124  
 FN\_ctx\_t(3XFN) 126  
 fn\_ctx\_unbind(3XFN) 129  
 FN\_identifier\_t(3XFN) 130  
 FN\_ref\_addr\_t(3XFN) 131  
 FN\_ref\_t(3XFN) 133  
 FN\_search\_control\_t(3XFN) 136  
 FN\_search\_filter\_t(3XFN) 139  
 FN\_status\_t(3XFN) 146  
 FN\_string\_t(3XFN) 151  
 getaddrinfo(3SOCKET) 155  
 gethostbyname(3NSL) 159

gethostname(3XNET)	165
getipnodebyname(3SOCKET)	166
getnetbyname(3SOCKET)	172
getnetconfig(3NSL)	176
getnetpath(3NSL)	178
getpeername(3SOCKET)	180
getpeername(3XNET)	181
getprotobyname(3SOCKET)	182
getpublickey(3NSL)	185
getrpcbyname(3NSL)	186
getservbyname(3SOCKET)	189
getsockname(3SOCKET)	193
getsockname(3XNET)	194
getsockopt(3SOCKET)	195
getsockopt(3XNET)	199
gss_accept_sec_context(3GSS)	202
gss_acquire_cred(3GSS)	208
gss_add_cred(3GSS)	211
gss_add_oid_set_member(3GSS)	215
gss_canonicalize_name(3GSS)	216
gss_compare_name(3GSS)	218
gss_context_time(3GSS)	219
gss_create_empty_oid_set(3GSS)	220
gss_delete_sec_context(3GSS)	221
gss_display_name(3GSS)	223
gss_display_status(3GSS)	225
gss_duplicate_name(3GSS)	227
gss_export_name(3GSS)	228
gss_export_sec_context(3GSS)	229
gss_get_mic(3GSS)	231
gss_import_name(3GSS)	233
gss_import_sec_context(3GSS)	235
gss_indicate_mechs(3GSS)	237
gss_init_sec_context(3GSS)	238
gss_inquire_context(3GSS)	245
gss_inquire_cred(3GSS)	248
gss_inquire_cred_by_mech(3GSS)	250
gss_inquire_mechs_for_name(3GSS)	252

gss\_inquire\_names\_for\_mech(3GSS) 254  
gss\_oid\_to\_str(3GSS) 255  
gss\_process\_context\_token(3GSS) 257  
gss\_release\_buffer(3GSS) 259  
gss\_release\_cred(3GSS) 260  
gss\_release\_name(3GSS) 261  
gss\_release\_oid(3GSS) 262  
gss\_release\_oid\_set(3GSS) 263  
gss\_str\_to\_oid(3GSS) 264  
gss\_test\_oid\_set\_member(3GSS) 266  
gss\_unwrap(3GSS) 267  
gss\_verify\_mic(3GSS) 269  
gss\_wrap(3GSS) 271  
gss\_wrap\_size\_limit(3GSS) 273  
htonl(3XNET) 275  
if\_nametoindex(3NSL) 276  
if\_nametoindex(3XNET) 278  
inet(3SOCKET) 280  
inet\_addr(3XNET) 284  
ldap(3LDAP) 286  
ldap\_abandon(3LDAP) 297  
ldap\_add(3LDAP) 298  
ldap\_ber\_free(3LDAP) 300  
ldap\_bind(3LDAP) 301  
ldap\_charset(3LDAP) 304  
ldap\_compare(3LDAP) 306  
ldap\_control\_free(3LDAP) 308  
ldap\_delete(3LDAP) 309  
ldap\_disptmpl(3LDAP) 310  
ldap\_entry2text(3LDAP) 316  
ldap\_error(3LDAP) 319  
ldap\_first\_attribute(3LDAP) 323  
ldap\_first\_entry(3LDAP) 324  
ldap\_first\_message(3LDAP) 326  
ldap\_friendly(3LDAP) 327  
ldap\_get\_dn(3LDAP) 328  
ldap\_get\_entry\_controls(3LDAP) 330  
ldap\_getfilter(3LDAP) 331

ldap_get_lang_values(3LDAP)	333
ldap_get_option(3LDAP)	335
ldap_get_values(3LDAP)	338
ldap_memcache(3LDAP)	340
ldap_memfree(3LDAP)	343
ldap_modify(3LDAP)	344
ldap_modrdn(3LDAP)	346
ldap_open(3LDAP)	348
ldap_parse_result(3LDAP)	350
ldap_result(3LDAP)	351
ldap_search(3LDAP)	353
ldap_searchprefs(3LDAP)	355
ldap_sort(3LDAP)	357
ldap_ufn(3LDAP)	359
ldap_url(3LDAP)	361
ldap_version(3LDAP)	364
listen(3SOCKET)	365
listen(3XNET)	366
netdir(3NSL)	368
nis_error(3NSL)	372
nis_groups(3NSL)	374
nis_local_names(3NSL)	377
nis_names(3NSL)	379
nis_objects(3NSL)	385
nis_ping(3NSL)	393
nis_server(3NSL)	394
nis_subr(3NSL)	396
nis_tables(3NSL)	399
nlsgetcall(3NSL)	408
nlsprovider(3NSL)	409
nlsrequest(3NSL)	410
rcmd(3SOCKET)	412
recv(3SOCKET)	414
recv(3XNET)	417
recvfrom(3XNET)	420
recvmsg(3XNET)	423
resolver(3RESOLV)	426
rexec(3SOCKET)	432

rpc(3NSL) 434  
rpcbind(3NSL) 443  
rpc\_clnt\_auth(3NSL) 445  
rpc\_clnt\_calls(3NSL) 447  
rpc\_clnt\_create(3NSL) 451  
rpc\_control(3NSL) 458  
rpc\_gss\_getcred(3NSL) 460  
rpc\_gss\_get\_error(3NSL) 462  
rpc\_gss\_get\_mechanisms(3NSL) 463  
rpc\_gss\_get\_principal\_name(3NSL) 465  
rpc\_gss\_max\_data\_length(3NSL) 467  
rpc\_gss\_mech\_to\_oid(3NSL) 468  
rpc\_gss\_seccreate(3NSL) 470  
rpc\_gss\_set\_callback(3NSL) 472  
rpc\_gss\_set\_defaults(3NSL) 474  
rpc\_gss\_set\_svc\_name(3NSL) 475  
rpc\_rac(3RAC) 477  
rpcsec\_gss(3NSL) 481  
rpc\_soc(3NSL) 486  
rpc\_svc\_calls(3NSL) 496  
rpc\_svc\_create(3NSL) 500  
rpc\_svc\_err(3NSL) 505  
rpc\_svc\_input(3NSL) 507  
rpc\_svc\_reg(3NSL) 509  
rpc\_xdr(3NSL) 511  
rstat(3RPC) 513  
rusers(3RPC) 514  
rwall(3RPC) 515  
secure\_rpc(3NSL) 516  
send(3SOCKET) 520  
send(3XNET) 522  
sendmsg(3XNET) 525  
sendto(3XNET) 529  
setsockopt(3XNET) 533  
shutdown(3SOCKET) 536  
shutdown(3XNET) 537  
slp\_api(3SLP) 538  
SLPclose(3SLP) 548



SLPDelAttrs(3SLP)	549
SLPDereg(3SLP)	551
SLPEscape(3SLP)	553
SLPFindAttrs(3SLP)	555
SLPFindScopes(3SLP)	557
SLPFindSrvs(3SLP)	559
SLPFindSrvTypes(3SLP)	561
SLPFree(3SLP)	563
SLPGetProperty(3SLP)	564
SLPGetRefreshInterval(3SLP)	565
SLPOpen(3SLP)	566
SLPParseSrvURL(3SLP)	568
SLPReg(3SLP)	570
SLPSetProperty(3SLP)	572
slp_strerror(3SLP)	573
SLPUnescape(3SLP)	574
socket(3SOCKET)	576
socket(3XNET)	579
socketpair(3SOCKET)	581
socketpair(3XNET)	582
spray(3SOCKET)	584
t_accept(3NSL)	586
t_alloc(3NSL)	590
t_bind(3NSL)	593
t_close(3NSL)	597
t_connect(3NSL)	599
t_errno(3NSL)	603
t_error(3NSL)	605
t_free(3NSL)	607
t_getinfo(3NSL)	609
t_getprotaddr(3NSL)	613
t_getstate(3NSL)	615
t_listen(3NSL)	617
t_look(3NSL)	620
t_open(3NSL)	622
t_optmgmt(3NSL)	626
t_rcv(3NSL)	634
t_rcvconnect(3NSL)	637

t\_rcvdis(3NSL) 640  
t\_rcvrel(3NSL) 642  
t\_rcvreldata(3NSL) 644  
t\_rcvudata(3NSL) 646  
t\_rcvuderr(3NSL) 649  
t\_rcvv(3NSL) 651  
t\_rcvvudata(3NSL) 654  
t\_snd(3NSL) 656  
t\_snddis(3NSL) 660  
t\_sndrel(3NSL) 662  
t\_sndreldata(3NSL) 664  
t\_sndudata(3NSL) 666  
t\_sndv(3NSL) 669  
t\_sndvudata(3NSL) 673  
t\_strerror(3NSL) 676  
t\_sync(3NSL) 677  
t\_sysconf(3NSL) 679  
t\_unbind(3NSL) 680  
xdr(3NSL) 682  
xdr\_admin(3NSL) 684  
xdr\_complex(3NSL) 686  
xdr\_create(3NSL) 689  
xdr\_simple(3NSL) 691  
xfn(3XFN) 695  
xfn\_attributes(3XFN) 696  
xfn\_composite\_names(3XFN) 699  
xfn\_compound\_names(3XFN) 700  
xfn\_links(3XFN) 703  
xfn\_status\_codes(3XFN) 706  
ypclnt(3NSL) 710  
yp\_update(3NSL) 715

**Index** 717

# Preface

---

Both novice users and those familiar with the SunOS operating system can use online man pages to obtain information about the system and its features. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

---

## Overview

The following contains a brief description of each man page section and the information it references:

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2.
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
- Section 5 contains miscellaneous documentation such as character-set tables.
- Section 6 contains available games and demos.
- Section 7 describes various special files that refer to specific hardware peripherals and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.

- Section 9 provides reference information needed to write device drivers in the kernel environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver/Kernel Interface (DKI).
- Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer can include in a device driver.
- Section 9F describes the kernel functions available for use by device drivers.
- Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

Below is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if there are no bugs to report, there is no BUGS section. See the `intro` pages for more information and detail about each section, and `man(1)` for more information about man pages in general.

NAME	This section gives the names of the commands or functions documented, followed by a brief description of what they do.
SYNOPSIS	This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full path name is shown. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.
	The following special characters are used in this section:
[ ]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.
. . .	Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .".
	Separator. Only one of the arguments separated by this character can be specified at a time.
{ }	Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.

PROTOCOL	This section occurs only in subsection 3R to indicate the protocol description file.
DESCRIPTION	This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES. Interactive commands, subcommands, requests, macros, and functions are described under USAGE.
IOCTL	This section appears on pages in Section 7 only. Only the device class that supplies appropriate parameters to the <code>ioctl(2)</code> system call is called <code>ioctl</code> and generates its own heading. <code>ioctl</code> calls for a specific device are listed alphabetically (on the man page for that specific device). <code>ioctl</code> calls are used for a particular class of devices all of which have an <code>io</code> ending, such as <code>mtio(7I)</code> .
OPTIONS	This section lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.
OPERANDS	This section lists the command operands and describes how they affect the actions of the command.
OUTPUT	This section describes the output – standard output, standard error, or output files – generated by the command.
RETURN VALUES	If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES.
ERRORS	On failure, most functions place an error code in the global variable <code>errno</code> indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than

	one condition can cause the same error, each condition is described in a separate paragraph under the error code.
USAGE	This section lists special rules, features, and commands that require in-depth explanations. The subsections listed here are used to explain built-in functionality:  Commands Modifiers Variables Expressions Input Grammar
EXAMPLES	This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command-line entry and machine response is shown. Whenever an example is given, the prompt is shown as <code>example%</code> , or if the user must be superuser, <code>example#</code> . Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS, and USAGE sections.
ENVIRONMENT VARIABLES	This section lists any environment variables that the command or function affects, followed by a brief description of the effect.
EXIT STATUS	This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion, and values other than zero for various error conditions.
FILES	This section lists all file names referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.
ATTRIBUTES	This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See <code>attributes(5)</code> for more information.
SEE ALSO	This section lists references to other man pages, in-house documentation, and outside publications.

DIAGNOSTICS	This section lists diagnostic messages with a brief explanation of the condition causing the error.
WARNINGS	This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.
NOTES	This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here.
BUGS	This section describes known bugs and, wherever possible, suggests workarounds.





# Networking Library Functions

---

## accept(3SOCKET)

<b>NAME</b>	accept – accept a connection on a socket										
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lsocket -lnsl [ <i>library</i> ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  int <b>accept</b>(int <i>s</i>, struct sockaddr *<i>addr</i>, socklen_t *<i>addrlen</i>);</pre>										
<b>DESCRIPTION</b>	<p>The argument <i>s</i> is a socket that has been created with <code>socket(3SOCKET)</code> and bound to an address with <code>bind(3SOCKET)</code>, and that is listening for connections after a call to <code>listen(3SOCKET)</code>. The <code>accept()</code> function extracts the first connection on the queue of pending connections, creates a new socket with the properties of <i>s</i>, and allocates a new file descriptor, <i>ns</i>, for the socket. If no pending connections are present on the queue and the socket is not marked as non-blocking, <code>accept()</code> blocks the caller until a connection is present. If the socket is marked as non-blocking and no pending connections are present on the queue, <code>accept()</code> returns an error as described below. The <code>accept()</code> function uses the <code>netconfig(4)</code> file to determine the STREAMS device file name associated with <i>s</i>. This is the device on which the connect indication will be accepted. The accepted socket, <i>ns</i>, is used to read and write data to and from the socket that connected to <i>ns</i>. It is not used to accept more connections. The original socket (<i>s</i>) remains open for accepting further connections.</p> <p>The argument <i>addr</i> is a result parameter that is filled in with the address of the connecting entity as it is known to the communications layer. The exact format of the <i>addr</i> parameter is determined by the domain in which the communication occurs.</p> <p>The argument <i>addrlen</i> is a value-result parameter. Initially, it contains the amount of space pointed to by <i>addr</i>; on return it contains the length in bytes of the address returned.</p> <p>The <code>accept()</code> function is used with connection-based socket types, currently with <code>SOCK_STREAM</code>.</p> <p>It is possible to <code>select(3C)</code> or <code>poll(2)</code> a socket for the purpose of an <code>accept()</code> by selecting or polling it for a read. However, this will only indicate when a connect indication is pending; it is still necessary to call <code>accept()</code>.</p>										
<b>RETURN VALUES</b>	The <code>accept()</code> function returns <code>-1</code> on error. If it succeeds, it returns a non-negative integer that is a descriptor for the accepted socket.										
<b>ERRORS</b>	<p><code>accept()</code> will fail if:</p> <table><tr><td>EBADF</td><td>The descriptor is invalid.</td></tr><tr><td>ECONNABORTED</td><td>The remote side aborted the connection before the <code>accept()</code> operation completed.</td></tr><tr><td>EFAULT</td><td>The <i>addr</i> parameter or the <i>addrlen</i> parameter is invalid.</td></tr><tr><td>EINTR</td><td>The <code>accept()</code> attempt was interrupted by the delivery of a signal.</td></tr><tr><td>EMFILE</td><td>The per-process descriptor table is full.</td></tr></table>	EBADF	The descriptor is invalid.	ECONNABORTED	The remote side aborted the connection before the <code>accept()</code> operation completed.	EFAULT	The <i>addr</i> parameter or the <i>addrlen</i> parameter is invalid.	EINTR	The <code>accept()</code> attempt was interrupted by the delivery of a signal.	EMFILE	The per-process descriptor table is full.
EBADF	The descriptor is invalid.										
ECONNABORTED	The remote side aborted the connection before the <code>accept()</code> operation completed.										
EFAULT	The <i>addr</i> parameter or the <i>addrlen</i> parameter is invalid.										
EINTR	The <code>accept()</code> attempt was interrupted by the delivery of a signal.										
EMFILE	The per-process descriptor table is full.										

accept(3SOCKET)

ENODEV	The protocol family and type corresponding to <i>s</i> could not be found in the <code>netconfig</code> file.
ENOMEM	There was insufficient user memory available to complete the operation.
ENOSR	There were insufficient STREAMS resources available to complete the operation.
ENOTSOCK	The descriptor does not reference a socket.
EOPNOTSUPP	The referenced socket is not of type <code>SOCK_STREAM</code> .
EPROTO	A protocol error has occurred; for example, the STREAMS protocol stack has not been initialized or the connection has already been released.
EWOULDBLOCK	The socket is marked as non-blocking and no connections are present to be accepted.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `poll(2)`, `bind(3SOCKET)`, `connect(3SOCKET)`, `listen(3SOCKET)`, `select(3C)`, `socket(3HEAD)`, `socket(3SOCKET)`, `netconfig(4)`, `attributes(5)`

## accept(3XNET)

<b>NAME</b>	accept – accept a new connection on a socket						
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  int <b>accept</b>(int <i>socket</i>, struct sockaddr *<i>address</i>, socklen_t            *<i>address_len</i>);</pre>						
<b>DESCRIPTION</b>	<p>The <code>accept()</code> function extracts the first connection on the queue of pending connections, creates a new socket with the same socket type protocol and address family as the specified socket, and allocates a new file descriptor for that socket.</p> <p>The function takes the following arguments:</p> <table><tr><td><i>socket</i></td><td>Specifies a socket that was created with <code>socket(3XNET)</code>, has been bound to an address with <code>bind(3XNET)</code>, and has issued a successful call to <code>listen(3XNET)</code>.</td></tr><tr><td><i>address</i></td><td>Either a null pointer, or a pointer to a <code>sockaddr</code> structure where the address of the connecting socket will be returned.</td></tr><tr><td><i>address_len</i></td><td>Points to a <code>socklen_t</code> which on input specifies the length of the supplied <code>sockaddr</code> structure, and on output specifies the length of the stored address.</td></tr></table> <p>If <i>address</i> is not a null pointer, the address of the peer for the accepted connection is stored in the <code>sockaddr</code> structure pointed to by <i>address</i>, and the length of this address is stored in the object pointed to by <i>address_len</i>.</p> <p>If the actual length of the address is greater than the length of the supplied <code>sockaddr</code> structure, the stored address will be truncated.</p> <p>If the protocol permits connections by unbound clients, and the peer is not bound, then the value stored in the object pointed to by <i>address</i> is unspecified.</p> <p>If the listen queue is empty of connection requests and <code>O_NONBLOCK</code> is not set on the file descriptor for the socket, <code>accept()</code> will block until a connection is present. If the <code>listen(3XNET)</code> queue is empty of connection requests and <code>O_NONBLOCK</code> is set on the file descriptor for the socket, <code>accept()</code> will fail and set <code>errno</code> to <code>EAGAIN</code> or <code>EWOULDBLOCK</code>.</p> <p>The accepted socket cannot itself accept more connections. The original socket remains open and can accept more connections.</p>	<i>socket</i>	Specifies a socket that was created with <code>socket(3XNET)</code> , has been bound to an address with <code>bind(3XNET)</code> , and has issued a successful call to <code>listen(3XNET)</code> .	<i>address</i>	Either a null pointer, or a pointer to a <code>sockaddr</code> structure where the address of the connecting socket will be returned.	<i>address_len</i>	Points to a <code>socklen_t</code> which on input specifies the length of the supplied <code>sockaddr</code> structure, and on output specifies the length of the stored address.
<i>socket</i>	Specifies a socket that was created with <code>socket(3XNET)</code> , has been bound to an address with <code>bind(3XNET)</code> , and has issued a successful call to <code>listen(3XNET)</code> .						
<i>address</i>	Either a null pointer, or a pointer to a <code>sockaddr</code> structure where the address of the connecting socket will be returned.						
<i>address_len</i>	Points to a <code>socklen_t</code> which on input specifies the length of the supplied <code>sockaddr</code> structure, and on output specifies the length of the stored address.						
<b>USAGE</b>	When a connection is available, <code>select(3C)</code> will indicate that the file descriptor for the socket is ready for reading.						
<b>RETURN VALUES</b>	Upon successful completion, <code>accept()</code> returns the nonnegative file descriptor of the accepted socket. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.						
<b>ERRORS</b>	The <code>accept()</code> function will fail if:						

EAGAIN	
EWOULDBLOCK	O_NONBLOCK is set for the socket file descriptor and no connections are present to be accepted.
EBADF	The <i>socket</i> argument is not a valid file descriptor.
ECONNABORTED	A connection has been aborted.
EFAULT	The <i>address</i> or <i>address_len</i> parameter can not be accessed or written.
EINTR	The <code>accept ()</code> function was interrupted by a signal that was caught before a valid connection arrived.
EINVAL	The <i>socket</i> is not accepting connections.
EMFILE	OPEN_MAX file descriptors are currently open in the calling process.
ENFILE	The maximum number of file descriptors in the system are already open.
ENOTSOCK	The <i>socket</i> argument does not refer to a socket.
EOPNOTSUPP	The socket type of the specified socket does not support accepting connections.

The `accept ()` function may fail if:

ENOBUFS	No buffer space is available.
ENOMEM	There was insufficient memory available to complete the operation.
ENOSR	There was insufficient STREAMS resources available to complete the operation.
EPROTO	A protocol error has occurred; for example, the STREAMS protocol stack has not been initialized.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `bind(3XNET)`, `connect(3XNET)`, `listen(3XNET)`, `socket(3XNET)`, `attributes(5)`

## ber\_decode(3LDAP)

<b>NAME</b>	ber_decode, ber_alloc_t, ber_free, ber_bvdup, ber_init, ber_flatten, ber_get_next, ber_skiptag, ber_peek_tag, ber_scanf, ber_get_int, ber_get_stringa, ber_get_stringal, ber_get_stringb, ber_get_null, ber_get_boolean, ber_get_bitstring, ber_first_element, ber_next_element, ber_bvfree, ber_bvecfree – Basic Encoding Rules library decoding functions
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt;  BerElement *ber_alloc_t(int options);  struct berval *ber_bvdup(struct berval *bv);  void ber_free(BerElement *ber, int freebuf);  BerElement *ber_init(struct berval *bv);  int ber_flatten(BerElement *ber, struct berval **bvPtr);  ber_get_next(Sockbuf *sb, unsigned long *len, char *bv_val);  ber_skip_tag(BerElement **ber, unsigned long **len);  ber_peek_tag(BerElement **ber, unsigned long **len);  ber_get_int(BerElement **ber, long **num);  ber_get_stringb(BerElement **ber, char **buf, unsigned long **len);  ber_get_stringa(BerElement **ber, char ***buf);  ber_get_stringal(BerElement **ber, struct berval ***bv);  ber_get_null(BerElement **ber);  ber_get_boolean(BerElement **ber, int **bool);  ber_get_bitstringa(BerElement **ber, char ***buf, unsigned long **blen);  ber_first_element(BerElement **ber, unsigned long **len, char ***cookie);  ber_next_element(BerElement **ber, unsigned long **len, char **cookie);  ber_scanf(BerElement **ber, char **fmt [, arg...]);  ber_bvfree(struct berval **bv);  ber_bvecfree(struct berval ***bvec);</pre>
<b>DESCRIPTION</b>	<p>These functions provide a subfunction interface to a simplified implementation of the Basic Encoding Rules of ASN.1. The version of BER these functions support is the one defined for the LDAP protocol. The encoding rules are the same as BER, except that only definite form lengths are used, and bitstrings and octet strings are always encoded in primitive form. In addition, these lightweight BER functions restrict tags and class to fit in a single octet (this means the actual tag must be less than 31). When</p>

a "tag" is specified in the descriptions below, it refers to the tag, class, and primitive or constructed bit in the first octet of the encoding. This man page describes the decoding functions in the lber library. See `ber_encode(3LDAP)` for details on the corresponding encoding functions.

Normally, the only functions that need be called by an application are `ber_get_next()` to get the next BER element and `ber_scanf()` to do the actual decoding. In some cases, `ber_peek_tag()` may also need to be called in normal usage. The other functions are provided for those applications that need more control than `ber_scanf()` provides. In general, these functions return the tag of the element decoded, or `-1` if an error occurred.

The `ber_get_next()` function is used to read the next BER element from the given Sockbuf, *sb*. A Sockbuf consists of the descriptor (usually socket, but a file descriptor works just as well) from which to read, and a BerElement structure used to maintain a buffer. On the first call, the *sb\_ber* struct should be zeroed. It strips off and returns the leading tag byte, strips off and returns the length of the entire element in *len*, and sets up *ber* for subsequent calls to `ber_scanf()`, and all to decode the element.

The `ber_scanf()` function is used to decode a BER element in much the same way that `scanf(3C)` works. It reads from *ber*, a pointer to a BerElement such as returned by `ber_get_next()`, interprets the bytes according to the format string *fmt*, and stores the results in its additional arguments. The format string contains conversion specifications which are used to direct the interpretation of the BER element. The format string can contain the following characters.

- a                    Octet string. A char \*\* should be supplied. Memory is allocated, filled with the contents of the octet string, null-terminated, and returned in the parameter.
- s                    Octet string. A char \* buffer should be supplied, followed by a pointer to an integer initialized to the size of the buffer. Upon return, the null-terminated octet string is put into the buffer, and the integer is set to the actual size of the octet string.
- O                    Octet string. A struct ber\_val \*\* should be supplied, which upon return points to a memory allocated struct berval containing the octet string and its length. `ber_bvfree()` can be called to free the allocated memory.
- b                    Boolean. A pointer to an integer should be supplied.
- i                    Integer. A pointer to an integer should be supplied.
- B                    Bitstring. A char \*\* should be supplied which will point to the memory allocated bits, followed by an unsigned long \*, which will point to the length (in bits) of the bitstring returned.
- n                    Null. No parameter is required. The element is simply skipped if it is recognized.

## ber\_decode(3LDAP)

- v Sequence of octet strings. A char \*\*\* should be supplied, which upon return points to a memory allocated null-terminated array of char \*'s containing the octet strings. NULL is returned if the sequence is empty.
- V Sequence of octet strings with lengths. A struct berval \*\*\* should be supplied, which upon return points to a memory allocated, null-terminated array of struct berval \*'s containing the octet strings and their lengths. NULL is returned if the sequence is empty. `ber_bvecfree()` can be called to free the allocated memory.
- x Skip element. The next element is skipped.
- { Begin sequence. No parameter is required. The initial sequence tag and length are skipped.
- } End sequence. No parameter is required and no action is taken.
- ] & Begin set. No parameter is required. The initial set tag and length are skipped.
- ] End set. No parameter is required and no action is taken.

The `ber_get_int()` function tries to interpret the next element as an integer, returning the result in *num*. The tag of whatever it finds is returned on success, -1 on failure.

The `ber_get_stringb()` function is used to read an octet string into a preallocated buffer. The *len* parameter should be initialized to the size of the buffer, and will contain the length of the octet string read upon return. The buffer should be big enough to take the octet string value plus a terminating NULL byte.

The `ber_get_stringa()` function is used to allocate memory space into which an octet string is read.

The `ber_get_stringal()` function is used to allocate memory space into which an octet string and its length are read. It takes a struct berval \*\*, and returns the result in this parameter.

The `ber_get_null()` function is used to read a NULL element. It returns the tag of the element it skips over.

The `ber_get_boolean()` function is used to read a boolean value. It is called the same way that `ber_get_int()` is called.

The `ber_get_bitstringa()` function is used to read a bitstring value. It takes a char \*\* which will hold the allocated memory bits, followed by an unsigned long \*, which will point to the length (in bits) of the bitstring returned.



The `ber_first_element()` function is used to return the tag and length of the first element in a set or sequence. It also returns in *cookie* a magic cookie parameter that should be passed to subsequent calls to `ber_next_element()`, which returns similar information.

`ber_alloc_t()` constructs and returns `BerElement`. A null pointer is returned on error. The options field contains a bitwise-or of options which are to be used when generating the encoding of this `BerElement`. One option is defined and must always be supplied:

```
#define LBER_USE_DER 0x01
```

When this option is present, lengths will always be encoded in the minimum number of octets. Note that this option does not cause values of sets and sequences to be rearranged in tag and byte order, so these functions are not suitable for generating DER output as defined in X.509 and X.680

The `ber_init` function constructs a `BerElement` and returns a new `BerElement` containing a copy of the data in the *bv* argument. `ber_init` returns the null pointer on error.

`ber_free()` frees a `BerElement` which is returned from the API calls `ber_alloc_t()` or `ber_init()`. Each `BerElement` must be freed by the caller. The second argument *freebuf* should always be set to 1 to ensure that the internal buffer used by the BER functions is freed as well as the `BerElement` container itself.

`ber_bvdup()` returns a copy of a *berval*. The *bv\_val* field in the returned *berval* points to a different area of memory as the *bv\_val* field in the argument *berval*. The null pointer is returned on error (that is, is out of memory).

The `ber_flatten` routine allocates a struct *berval* whose contents are BER encoding taken from the *ber* argument. The *bvPtr* pointer points to the returned *berval*, which must be freed using `ber_bvfree()`. This routine returns 0 on success and -1 on error.

## EXAMPLES

**EXAMPLE 1** Assume the variable *ber* contains a lightweight BER encoding of the following ASN.1 object:

```
AlmostASearchRequest := SEQUENCE {
    baseObject      DistinguishedName,
    scope           ENUMERATED {
        baseObject      (0),
        singleLevel     (1),
        wholeSubtree    (2)
    },
    derefAliases    ENUMERATED {
        neverDerefaliases (0),
        derefInSearching  (1),
        derefFindingBaseObj (2),
        alwaysDerefAliases (3N)
    },
    sizelimit       INTEGER (0 .. 65535),
    timelimit       INTEGER (0 .. 65535),
```

## ber\_decode(3LDAP)

**EXAMPLE 1** Assume the variable *ber* contains a lightweight BER encoding of the following ASN.1 object: (Continued)

```
        attrsOnly      BOOLEAN,
        attributes     SEQUENCE OF AttributeType
    }
```

**EXAMPLE 2** The element can be decoded using `ber_scanf()` as follows.

```
int    scope, ali, size, time, attrsonly;
char  *dn, **attrs;
if ( ber_scanf( ber, "{aiiiib{v}}", &dn, &scope, &ali,
              &size, &time, &attrsonly, &attrs ) == -1 )
    /* error */
else
    /* success */
```

**ERRORS** If an error occurs during decoding, generally these functions return `-1`.

**NOTES** The return values for all of these functions are declared in the `<ldap.h>` header file. Some functions may allocate memory which must be freed by the calling application.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ber_encode(3LDAP)`

Yeong, W., Howes, T., and Hardcastle-Kille, S., "Lightweight Directory Access Protocol", OSI-DS-26, April 1992.

Information Processing - Open Systems Interconnection - Model and Notation - Service Definition - Specification of Basic Encoding Rules for Abstract Syntax Notation One, International Organization for Standardization, International Standard 8825.

<b>NAME</b>	ber_encode, ber_alloc, ber_printf, ber_put_int, ber_put_ostring, ber_put_string, ber_put_null, ber_put_boolean, ber_put_bitstring, ber_start_seq, ber_start_set, ber_put_seq, ber_put_set – simplified Basic Encoding Rules library encoding functions
<b>SYNOPSIS</b>	<pre>cc[ <i>flag...</i> ] <i>file...</i> -lldap[ <i>library...</i> ]  #include &lt;lber.h&gt;  BerElement*ber_alloc();  ber_printf(BerElement *ber, char **fmt[, arg...]);  ber_put_int(BerElement *ber, long num, char tag);  ber_put_ostring(BerElement *ber, char **str, unsigned long len, char tag);  ber_put_string(BerElement *ber, char **str, char tag);  ber_put_null(BerElement *ber, char tag);  ber_put_boolean(BerElement *ber, int bool, char tag);  ber_put_bitstring(BerElement *ber, char *str, int blen, char tag);  ber_start_seq(BerElement *ber, char tag);  ber_start_set(BerElement *ber, char tag);  ber_put_seq(BerElement *ber);  ber_put_set(BerElement *ber);</pre>
<b>DESCRIPTION</b>	<p>These functions provide a subfunction interface to a simplified implementation of the Basic Encoding Rules of ASN.1. The version of BER these functions support is the one defined for the LDAP protocol. The encoding rules are the same as BER, except that only definite form lengths are used, and bitstrings and octet strings are always encoded in primitive form. In addition, these lightweight BER functions restrict tags and class to fit in a single octet (this means the actual tag must be less than 31). When a "tag" is specified in the descriptions below, it refers to the tag, class, and primitive or constructed bit in the first octet of the encoding. This man page describes the encoding functions in the lber library. See <code>ber_decode(3LDAP)</code> for details on the corresponding decoding functions.</p> <p>Normally, the only functions that need be called by an application are <code>ber_alloc()</code>, to allocate a BER element, and <code>ber_printf()</code> to do the actual encoding. The other functions are provided for those applications that need more control than <code>ber_printf()</code> provides. In general, these functions return the length of the element encoded, or -1 if an error occurred.</p> <p>The <code>ber_alloc()</code> function is used to allocate a new BER element.</p> <p>The <code>ber_printf()</code> function is used to encode a BER element in much the same way that <code>sprintf(3S)</code> works. One important difference, though, is that some state information is kept with the <code>ber</code> parameter so that multiple calls can be made to</p>

## ber\_encode(3LDAP)

`ber_printf()` to append things to the end of the BER element. `Ber_printf()` writes to *ber*, a pointer to a `BerElement` such as returned by `ber_alloc()`. It interprets and formats its arguments according to the format string *fmt*. The format string can contain the following characters:

-b	Boolean. An integer parameter should be supplied. A boolean element is output.
-i	Integer. An integer parameter should be supplied. An integer element is output.
-B	Bitstring. A <code>char *</code> pointer to the start of the bitstring is supplied, followed by the number of bits in the bitstring. A bitstring element is output.
-n	Null. No parameter is required. A null element is output.
-o	Octet string. A <code>char *</code> is supplied, followed by the length of the string pointed to. An octet string element is output.
-s	Octet string. A null-terminated string is supplied. An octet string element is output, not including the trailing NULL octet.
-t	Tag. An int specifying the tag to give the next element is provided. This works across calls.
-v	Several octet strings. A null-terminated array of <code>char *</code> s is supplied. Note that a construct like <code>'{v}'</code> is required to get an actual SEQUENCE OF octet strings.
-{	Begin sequence. No parameter is required.
-}	End sequence. No parameter is required.
- ] &	Begin set. No parameter is required.
-]	End set. No parameter is required.

The `ber_put_int()` function writes the integer element *num* to the BER element *ber*.

The `ber_put_boolean()` function writes the boolean value given by *bool* to the BER element.

The `ber_put_bitstring()` function writes *blen* bits starting at *str* as a bitstring value to the given BER element. Note that *blen* is the length in *bits* of the bitstring.

The `ber_put_ostring()` function writes *len* bytes starting at *str* to the BER element as an octet string.

The `ber_put_string()` function writes the null-terminated string (minus the terminating `"`) to the BER element as an octet string.

The `ber_put_null()` function writes a NULL element to the BER element.

The `ber_start_seq()` function is used to start a sequence in the BER element. The `ber_start_set()` function works similarly. The end of the sequence or set is marked by the nearest matching call to `ber_put_seq()` or `ber_put_set()`, respectively.

The `ber_first_element()` function is used to return the tag and length of the first element in a set or sequence. It also returns in *cookie* a magic cookie parameter that should be passed to subsequent calls to `ber_next_element()`, which returns similar information.

**EXAMPLES**

**EXAMPLE 1** Assuming the following variable declarations, and that the variables have been assigned appropriately, an BER encoding of the following ASN.1 object:

```

AlmostASearchRequest := SEQUENCE {
    baseObject      DistinguishedName,
    scope           ENUMERATED {
        baseObject      (0),
        singleLevel     (1),
        wholeSubtree    (2)
    },
    derefAliases    ENUMERATED {
        neverDerefaliases (0),
        derefInSearching  (1),
        derefFindingBaseObj (2),
        alwaysDerefAliases (3N)
    },
    sizelimit       INTEGER (0 .. 65535),
    timelimit       INTEGER (0 .. 65535),
    attrsOnly       BOOLEAN,
    attributes      SEQUENCE OF AttributeType
}

```

can be achieved like so:

```

int    scope, ali, size, time, attrsonly;
char   *dn, **attrs;

/* ... fill in values ... */
if ( (ber = ber_alloc( )) == NULLBER )
/* error */

if ( ber_printf( ber, "{siiiib{v}}", dn, scope, ali,
    size, time, attrsonly, attrs ) == -1 )
/* error */
else
/* success */

```

**RETURN VALUES**

If an error occurs during encoding, `ber_alloc()` returns NULL; other functions generally return -1.

**ATTRIBUTES**

See `attributes(5)` for a description of the following attributes:

ber\_encode(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `attributes(5)`, `ber_decode(3LDAP)`

Yeong, W., Howes, T., and Hardcastle-Kille, S., "Lightweight Directory Access Protocol", OSI-DS-26, April 1992.

Information Processing - Open Systems Interconnection - Model and Notation - Service Definition - Specification of Basic Encoding Rules for Abstract Syntax Notation One, International Organization for Standardization, International Standard 8825.

**NOTES** The return values for all of these functions are declared in the `<liber.h>` header file.

<b>NAME</b>	bind – bind a name to a socket																												
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lsocket -lnsl [ <i>library</i> ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  int <b>bind</b>(int <i>s</i>, const struct sockaddr *<i>name</i>, int <i>namelen</i>);</pre>																												
<b>DESCRIPTION</b>	bind() assigns a name to an unnamed socket. When a socket is created with socket(3SOCKET), it exists in a name space (address family) but has no name assigned. bind() requests that the name pointed to by <i>name</i> be assigned to the socket.																												
<b>RETURN VALUES</b>	If the bind is successful, 0 is returned. A return value of -1 indicates an error, which is further specified in the global errno.																												
<b>ERRORS</b>	<p>The bind() call will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EACCES</td> <td>The requested address is protected and the current user has inadequate permission to access it.</td> </tr> <tr> <td style="vertical-align: top;">EADDRINUSE</td> <td>The specified address is already in use.</td> </tr> <tr> <td style="vertical-align: top;">EADDRNOTAVAIL</td> <td>The specified address is not available on the local machine.</td> </tr> <tr> <td style="vertical-align: top;">EBADF</td> <td><i>s</i> is not a valid descriptor.</td> </tr> <tr> <td style="vertical-align: top;">EINVAL</td> <td><i>namelen</i> is not the size of a valid address for the specified address family.</td> </tr> <tr> <td style="vertical-align: top;">EINVAL</td> <td>The socket is already bound to an address.</td> </tr> <tr> <td style="vertical-align: top;">ENOSR</td> <td>There were insufficient STREAMS resources for the operation to complete.</td> </tr> <tr> <td style="vertical-align: top;">ENOTSOCK</td> <td><i>s</i> is a descriptor for a file, not a socket.</td> </tr> </table> <p>The following errors are specific to binding names in the UNIX domain:</p> <table border="0"> <tr> <td style="vertical-align: top;">EACCES</td> <td>Search permission is denied for a component of the path prefix of the pathname in <i>name</i>.</td> </tr> <tr> <td style="vertical-align: top;">EIO</td> <td>An I/O error occurred while making the directory entry or allocating the inode.</td> </tr> <tr> <td style="vertical-align: top;">EISDIR</td> <td>A null pathname was specified.</td> </tr> <tr> <td style="vertical-align: top;">ELOOP</td> <td>Too many symbolic links were encountered in translating the pathname in <i>name</i>.</td> </tr> <tr> <td style="vertical-align: top;">ENOENT</td> <td>A component of the path prefix of the pathname in <i>name</i> does not exist.</td> </tr> <tr> <td style="vertical-align: top;">ENOTDIR</td> <td>A component of the path prefix of the pathname in <i>name</i> is not a directory.</td> </tr> </table>	EACCES	The requested address is protected and the current user has inadequate permission to access it.	EADDRINUSE	The specified address is already in use.	EADDRNOTAVAIL	The specified address is not available on the local machine.	EBADF	<i>s</i> is not a valid descriptor.	EINVAL	<i>namelen</i> is not the size of a valid address for the specified address family.	EINVAL	The socket is already bound to an address.	ENOSR	There were insufficient STREAMS resources for the operation to complete.	ENOTSOCK	<i>s</i> is a descriptor for a file, not a socket.	EACCES	Search permission is denied for a component of the path prefix of the pathname in <i>name</i> .	EIO	An I/O error occurred while making the directory entry or allocating the inode.	EISDIR	A null pathname was specified.	ELOOP	Too many symbolic links were encountered in translating the pathname in <i>name</i> .	ENOENT	A component of the path prefix of the pathname in <i>name</i> does not exist.	ENOTDIR	A component of the path prefix of the pathname in <i>name</i> is not a directory.
EACCES	The requested address is protected and the current user has inadequate permission to access it.																												
EADDRINUSE	The specified address is already in use.																												
EADDRNOTAVAIL	The specified address is not available on the local machine.																												
EBADF	<i>s</i> is not a valid descriptor.																												
EINVAL	<i>namelen</i> is not the size of a valid address for the specified address family.																												
EINVAL	The socket is already bound to an address.																												
ENOSR	There were insufficient STREAMS resources for the operation to complete.																												
ENOTSOCK	<i>s</i> is a descriptor for a file, not a socket.																												
EACCES	Search permission is denied for a component of the path prefix of the pathname in <i>name</i> .																												
EIO	An I/O error occurred while making the directory entry or allocating the inode.																												
EISDIR	A null pathname was specified.																												
ELOOP	Too many symbolic links were encountered in translating the pathname in <i>name</i> .																												
ENOENT	A component of the path prefix of the pathname in <i>name</i> does not exist.																												
ENOTDIR	A component of the path prefix of the pathname in <i>name</i> is not a directory.																												

bind(3SOCKET)

EROFS The inode would reside on a read-only file system.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `unlink(2)`, `socket(3SOCKET)`, `attributes(5)`, `socket(3HEAD)`

**NOTES** 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 (using `unlink(2)`).

The rules used in name binding vary between communication domains.



<b>NAME</b>	bind – bind a name to a socket																
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  int <b>bind</b>(int <i>socket</i>, const struct sockaddr *<i>address</i>, socklen_t         <i>address_len</i>);</pre>																
<b>DESCRIPTION</b>	<p>The <code>bind()</code> function assigns an <i>address</i> to an unnamed socket. Sockets created with <code>socket(3XNET)</code> function are initially unnamed; they are identified only by their address family.</p> <p>The function takes the following arguments:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><i>socket</i></td> <td>Specifies the file descriptor of the socket to be bound.</td> </tr> <tr> <td><i>address</i></td> <td>Points to a <code>sockaddr</code> structure containing the address to be bound to the socket. The length and format of the address depend on the address family of the socket.</td> </tr> <tr> <td><i>address_len</i></td> <td>Specifies the length of the <code>sockaddr</code> structure pointed to by the <i>address</i> argument.</td> </tr> </table> <p>The socket in use may require the process to have appropriate privileges to use the <code>bind()</code> function.</p>	<i>socket</i>	Specifies the file descriptor of the socket to be bound.	<i>address</i>	Points to a <code>sockaddr</code> structure containing the address to be bound to the socket. The length and format of the address depend on the address family of the socket.	<i>address_len</i>	Specifies the length of the <code>sockaddr</code> structure pointed to by the <i>address</i> argument.										
<i>socket</i>	Specifies the file descriptor of the socket to be bound.																
<i>address</i>	Points to a <code>sockaddr</code> structure containing the address to be bound to the socket. The length and format of the address depend on the address family of the socket.																
<i>address_len</i>	Specifies the length of the <code>sockaddr</code> structure pointed to by the <i>address</i> argument.																
<b>USAGE</b>	An application program can retrieve the assigned socket name with the <code>getsockname(3XNET)</code> function.																
<b>RETURN VALUES</b>	Upon successful completion, <code>bind()</code> returns 0. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.																
<b>ERRORS</b>	<p>The <code>bind()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EADDRINUSE</td> <td>The specified address is already in use.</td> </tr> <tr> <td>EADDRNOTAVAIL</td> <td>The specified address is not available from the local machine.</td> </tr> <tr> <td>EAFNOSUPPORT</td> <td>The specified address is not a valid address for the address family of the specified socket.</td> </tr> <tr> <td>EBADF</td> <td>The <i>socket</i> argument is not a valid file descriptor.</td> </tr> <tr> <td>EFAULT</td> <td>The <i>address</i> argument can not be accessed.</td> </tr> <tr> <td>EINVAL</td> <td>The socket is already bound to an address, and the protocol does not support binding to a new address; or the socket has been shut down.</td> </tr> <tr> <td>ENOTSOCK</td> <td>The <i>socket</i> argument does not refer to a socket.</td> </tr> <tr> <td>EOPNOTSUPP</td> <td>The socket type of the specified socket does not support binding to an address.</td> </tr> </table>	EADDRINUSE	The specified address is already in use.	EADDRNOTAVAIL	The specified address is not available from the local machine.	EAFNOSUPPORT	The specified address is not a valid address for the address family of the specified socket.	EBADF	The <i>socket</i> argument is not a valid file descriptor.	EFAULT	The <i>address</i> argument can not be accessed.	EINVAL	The socket is already bound to an address, and the protocol does not support binding to a new address; or the socket has been shut down.	ENOTSOCK	The <i>socket</i> argument does not refer to a socket.	EOPNOTSUPP	The socket type of the specified socket does not support binding to an address.
EADDRINUSE	The specified address is already in use.																
EADDRNOTAVAIL	The specified address is not available from the local machine.																
EAFNOSUPPORT	The specified address is not a valid address for the address family of the specified socket.																
EBADF	The <i>socket</i> argument is not a valid file descriptor.																
EFAULT	The <i>address</i> argument can not be accessed.																
EINVAL	The socket is already bound to an address, and the protocol does not support binding to a new address; or the socket has been shut down.																
ENOTSOCK	The <i>socket</i> argument does not refer to a socket.																
EOPNOTSUPP	The socket type of the specified socket does not support binding to an address.																

## bind(3XNET)

If the address family of the socket is AF\_UNIX, then bind() will fail if:

EACCES	A component of the path prefix denies search permission, or the requested name requires writing in a directory with a mode that denies write permission.
EDESTADDRREQ	
EISDIR	The <i>address</i> argument is a null pointer.
EIO	An I/O error occurred.
ELOOP	Too many symbolic links were encountered in translating the pathname in <i>address</i> .
ENAMETOOLONG	A component of a pathname exceeded NAME_MAX characters, or an entire pathname exceeded PATH_MAX characters.
ENOENT	A component of the pathname does not name an existing file or the pathname is an empty string.
ENOTDIR	A component of the path prefix of the pathname in <i>address</i> is not a directory.
EROFS	The name would reside on a read-only filesystem.

The bind() function may fail if:

EACCES	The specified address is protected and the current user does not have permission to bind to it.
EINVAL	The <i>address_len</i> argument is not a valid length for the address family.
EISCONN	The socket is already connected.
ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds PATH_MAX.
ENOBUFS	Insufficient resources were available to complete the call.
ENOSR	There were insufficient STREAMS resources for the operation to complete.

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** connect(3XNET), getsockname(3XNET), listen(3XNET), socket(3XNET), attributes(5)

**NAME** byteorder, htonl, htons, ntohl, ntohs – convert values between host and network byte order

**SYNOPSIS**

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

uint32_t  htonl (uint32_t  hostlong) ;
uint16_t  htons (uint16_t  hostshort) ;
uint32_t  ntohl (uint32_t  netlong) ;
uint16_t  ntohs (uint16_t  netshort) ;
```

**DESCRIPTION**

These routines convert 16 and 32 bit quantities between network byte order and host byte order. On some architectures these routines are defined as NULL macros in the include file <netinet/in.h>. On other architectures, if their host byte order is different from network byte order, these routines are functional.

These routines are most often used in conjunction with Internet addresses and ports as returned by `gethostent()` and `getservent()`. See `gethostbyname(3NSL)` and `getservbyname(3SOCKET)`.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `gethostbyname(3NSL)`, `getservbyname(3SOCKET)`, `attributes(5)`, `inet(3HEAD)`

## cldap\_close(3LDAP)

- NAME** cldap\_close – dispose of connectionless LDAP pointer
- SYNOPSIS**

```
cc [ flag... ] file... -lldap [ library... ]  
  
#include <lber.h>  
#include <ldap.h>  
  
void cldap_close(LDAP *ld);
```
- DESCRIPTION** The `cldap_close()` function disposes of memory allocated by `cldap_open(3LDAP)`. It should be called when all CLDAP communication is complete.
- PARAMETERS** *ld* The LDAP pointer returned by a previous call to `cldap_open(3LDAP)`.
- ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

- SEE ALSO** `ldap(3LDAP)`, `cldap_open(3LDAP)`, `cldap_search_s(3LDAP)`, `cldap_setretryinfo(3LDAP)`

**NAME** cldap\_open – LDAP connectionless communication preparation

**SYNOPSIS**

```
cc [ flag... ] file... -lldap [ library... ]

#include <lber.h>
#include <ldap.h>

LDAP *cldap_open (char *host, int port);
```

**PARAMETERS** *host* The name of the host on which the LDAP server is running.  
*port* The port number to connect.

**DESCRIPTION** The cldap\_open() function is called to prepare for connectionless LDAP communication (over udp(7P)). It allocates an LDAP structure which is passed to future search requests.

If the default IANA-assigned port of 389 is desired, LDAP\_PORT should be specified for *port*. *host* can contain a space-separated list of hosts or addresses to try. cldap\_open() returns a pointer to an LDAP structure, which should be passed to subsequent calls to cldap\_search\_s(3LDAP), cldap\_setretryinfo(3LDAP), and cldap\_close(3LDAP). Certain fields in the LDAP structure can be set to indicate size limit, time limit, and how aliases are handled during operations. See ldap\_open(3LDAP) and <ldap.h> for more details.

**ERRORS** If an error occurs, cldap\_open() will return NULL and errno will be set appropriately.

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP) cldap\_search\_s(3LDAP), cldap\_setretryinfo(3LDAP), cldap\_close(3LDAP), attributes(5), udp(7P)

## cldap\_search\_s(3LDAP)

<b>NAME</b>	cldap_search_s – connectionless LDAP search
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int cldap_search_s(LDAP *ld, char *base, int scope, char *filter, char     *attrs, int attrsonly, LDAPMessage **res, char *logdn);</pre>
<b>DESCRIPTION</b>	<p>The <code>cldap_search_s()</code> function performs an LDAP search using the Connectionless LDAP (CLDAP) protocol.</p> <p><code>cldap_search_s()</code> has parameters and behavior identical to that of <code>ldap_search_s(3LDAP)</code>, except for the addition of the <code>logdn</code> parameter. <code>logdn</code> should contain a distinguished name to be used only for logging purposed by the LDAP server. It should be in the text format described by RFC 1779, <i>A String Representation of Distinguished Names</i>.</p>
<b>Retransmission Algorithm</b>	<p><code>cldap_search_s()</code> operates using the CLDAP protocol over <code>udp(7P)</code>. Since UDP is a non-reliable protocol, a retry mechanism is used to increase reliability. The <code>cldap_setretryinfo(3LDAP)</code> function can be used to set two retry parameters: <code>tries</code>, a count of the number of times to send a search request and <code>timeout</code>, an initial timeout that determines how long to wait for a response before re-trying. <code>timeout</code> is specified seconds. These values are stored in the <code>ld_cldaptries</code> and <code>ld_cldaptimeout</code> members of the <code>ld</code> LDAP structure, and the default values set in <code>ldap_open(3LDAP)</code> are 4 and 3 respectively. The retransmission algorithm used is:</p> <p>Step 1: Set the current timeout to <code>ld_cldaptimeout</code> seconds, and the current LDAP server address to the first LDAP server found during the <code>ldap_open(3LDAP)</code> call.</p> <p>Step 2: Send the search request to the current LDAP server address.</p> <p>Step 3: Set the wait timeout to the current timeout divided by the number of server addresses found during <code>ldap_open(3LDAP)</code> or to one second, whichever is larger. Wait at most that long for a response; if a response is received, STOP. Note that the wait timeout is always rounded down to the next lowest second.</p> <p>Step 5: Repeat steps 2 and 3 for each LDAP server address.</p> <p>Step 6: Set the current timeout to twice its previous value and repeat Steps 2 through 6 a maximum of <code>tries</code> times.</p>
<b>EXAMPLES</b>	<p>Assume that the default values for <code>tries</code> and <code>timeout</code> of 4 tries and 3 seconds are used. Further, assume that a space-separated list of two hosts, each with one address, was passed to <code>cldap_open(3LDAP)</code>. The pattern of requests sent will be (stopping as soon as a response is received):</p> <pre>Time          Search Request Sent To: +0           Host A try 1 +1 (0+3/2)    Host B try 1</pre>

## cldap\_search\_s(3LDAP)

```
+2 (1+3/2)      Host A try 2
+5 (2+6/2)      Host B try 2
+8 (5+6/2)      Host A try 3
+14 (8+12/2)    Host B try 3
+20 (14+12/2)   Host A try 4
+32 (20+24/2)   Host B try 4
+44 (20+24/2)   (give up - no response)
```

**ERRORS** cldap\_search\_s() returns LDAP\_SUCCESS if a search was successful and the appropriate LDAP error code otherwise. See ldap\_error(3LDAP) for more information.

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_error(3LDAP), ldap\_search\_s(3LDAP), cldap\_open(3LDAP), cldap\_setretryinfo(3LDAP), cldap\_close(3LDAP), attributes(5), udp(7P)

## cldap\_setretryinfo(3LDAP)

**NAME** cldap\_setretryinfo – set connectionless LDAP request retransmission parameters

**SYNOPSIS** `cc [ flag... ] file... -lldap [ library... ]`

```
#include <lber.h>
#include <ldap.h>

void cldap_setretryinfo(LDAP *ld, int tries, int timeout);
```

**PARAMETERS**

*ld* LDAP pointer returned from a previous call to cldap\_open(3LDAP).

*tries* Maximum number of times to send a request.

*timeout* Initial time, in seconds, to wait before re-sending a request.

**DESCRIPTION** The cldap\_setretryinfo() function is used to set the CLDAP request retransmission behavior for future cldap\_search\_s(3LDAP) calls. The default values (set by cldap\_open(3LDAP)) are 4 tries and 3 seconds between tries. See cldap\_search\_s(3LDAP) for a complete description of the retransmission algorithm used.

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), cldap\_open(3LDAP), cldap\_search\_s(3LDAP), cldap\_close(3LDAP), attributes(5)



<b>NAME</b>	connect – initiate a connection on a socket																		
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  int connect(int s, const struct sockaddr *name, int namelen);</pre>																		
<b>DESCRIPTION</b>	<p>The parameter <i>s</i> is a socket. If it is of type <code>SOCK_DGRAM</code>, <code>connect()</code> specifies the peer with which the socket is to be associated; this address is the address to which datagrams are to be sent if a receiver is not explicitly designated; it is the only address from which datagrams are to be received. If the socket <i>s</i> is of type <code>SOCK_STREAM</code>, <code>connect()</code> attempts to make a connection to another socket. The other socket is specified by <i>name</i>. <i>name</i> is an address in the communication space of the socket. Each communication space interprets the <i>name</i> parameter in its own way. If <i>s</i> is not bound, then it will be bound to an address selected by the underlying transport provider. Generally, stream sockets may successfully <code>connect()</code> only once; datagram sockets may use <code>connect()</code> multiple times to change their association. Datagram sockets may dissolve the association by connecting to a null address.</p>																		
<b>RETURN VALUES</b>	If the connection or binding succeeds, 0 is returned. Otherwise, -1 is returned and sets <code>errno</code> to indicate the error.																		
<b>ERRORS</b>	<p>The call fails if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EACCES</td> <td>Search permission is denied for a component of the path prefix of the pathname in <i>name</i>.</td> </tr> <tr> <td style="vertical-align: top;">EADDRINUSE</td> <td>The address is already in use.</td> </tr> <tr> <td style="vertical-align: top;">EADDRNOTAVAIL</td> <td>The specified address is not available on the remote machine.</td> </tr> <tr> <td style="vertical-align: top;">EAFNOSUPPORT</td> <td>Addresses in the specified address family cannot be used with this socket.</td> </tr> <tr> <td style="vertical-align: top;">EALREADY</td> <td>The socket is non-blocking and a previous connection attempt has not yet been completed.</td> </tr> <tr> <td style="vertical-align: top;">EBADF</td> <td><i>s</i> is not a valid descriptor.</td> </tr> <tr> <td style="vertical-align: top;">ECONNREFUSED</td> <td>The attempt to connect was forcefully rejected. The calling program should <code>close(2)</code> the socket descriptor, and issue another <code>socket(3SOCKET)</code> call to obtain a new descriptor before attempting another <code>connect()</code> call.</td> </tr> <tr> <td style="vertical-align: top;">EINPROGRESS</td> <td>The socket is non-blocking and the connection cannot be completed immediately. It is possible to <code>select(3C)</code> for completion by selecting the socket for writing.</td> </tr> <tr> <td style="vertical-align: top;">EINTR</td> <td>The connection attempt was interrupted before any data arrived by the delivery of a signal.</td> </tr> </table>	EACCES	Search permission is denied for a component of the path prefix of the pathname in <i>name</i> .	EADDRINUSE	The address is already in use.	EADDRNOTAVAIL	The specified address is not available on the remote machine.	EAFNOSUPPORT	Addresses in the specified address family cannot be used with this socket.	EALREADY	The socket is non-blocking and a previous connection attempt has not yet been completed.	EBADF	<i>s</i> is not a valid descriptor.	ECONNREFUSED	The attempt to connect was forcefully rejected. The calling program should <code>close(2)</code> the socket descriptor, and issue another <code>socket(3SOCKET)</code> call to obtain a new descriptor before attempting another <code>connect()</code> call.	EINPROGRESS	The socket is non-blocking and the connection cannot be completed immediately. It is possible to <code>select(3C)</code> for completion by selecting the socket for writing.	EINTR	The connection attempt was interrupted before any data arrived by the delivery of a signal.
EACCES	Search permission is denied for a component of the path prefix of the pathname in <i>name</i> .																		
EADDRINUSE	The address is already in use.																		
EADDRNOTAVAIL	The specified address is not available on the remote machine.																		
EAFNOSUPPORT	Addresses in the specified address family cannot be used with this socket.																		
EALREADY	The socket is non-blocking and a previous connection attempt has not yet been completed.																		
EBADF	<i>s</i> is not a valid descriptor.																		
ECONNREFUSED	The attempt to connect was forcefully rejected. The calling program should <code>close(2)</code> the socket descriptor, and issue another <code>socket(3SOCKET)</code> call to obtain a new descriptor before attempting another <code>connect()</code> call.																		
EINPROGRESS	The socket is non-blocking and the connection cannot be completed immediately. It is possible to <code>select(3C)</code> for completion by selecting the socket for writing.																		
EINTR	The connection attempt was interrupted before any data arrived by the delivery of a signal.																		

## connect(3SOCKET)

EINVAL	<i>namelen</i> is not the size of a valid address for the specified address family.
EIO	An I/O error occurred while reading from or writing to the file system.
EISCONN	The socket is already connected.
ELOOP	Too many symbolic links were encountered in translating the pathname in <i>name</i> .
ENETUNREACH	The network is not reachable from this host.
ENOENT	A component of the path prefix of the pathname in <i>name</i> does not exist.
ENOENT	The socket referred to by the pathname in <i>name</i> does not exist.
ENOSR	There were insufficient STREAMS resources available to complete the operation.
ENXIO	The server exited before the connection was complete.
ETIMEDOUT	Connection establishment timed out without establishing a connection.
EWOULDBLOCK	The socket is marked as non-blocking, and the requested operation would block.

The following errors are specific to connecting names in the UNIX domain. These errors may not apply in future versions of the UNIX IPC domain.

ENOTDIR	A component of the path prefix of the pathname in <i>name</i> is not a directory.
ENOTSOCK	<i>s</i> is not a socket.
ENOTSOCK	<i>name</i> is not a socket.
EPROTOTYPE	The file referred to by <i>name</i> is a socket of a type other than type <i>s</i> (for example, <i>s</i> is a SOCK_DGRAM socket, while <i>name</i> refers to a SOCK_STREAM socket).

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `close(2)`, `accept(3SOCKET)`, `getsockname(3SOCKET)`, `select(3C)`, `socket(3SOCKET)`, `socket(3HEAD)`, `attributes(5)`

<b>NAME</b>	connect – connect a socket
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  int <b>connect</b>(int <i>socket</i>, const struct sockaddr *<i>address</i>, socklen_t             <i>address_len</i>);</pre>
<b>DESCRIPTION</b>	<p>The <code>connect()</code> function requests a connection to be made on a socket. The function takes the following arguments:</p> <p><i>socket</i>                Specifies the file descriptor associated with the socket.</p> <p><i>address</i>               Points to a <code>sockaddr</code> structure containing the peer address. The length and format of the address depend on the address family of the socket.</p> <p><i>address_len</i>           Specifies the length of the <code>sockaddr</code> structure pointed to by the <i>address</i> argument.</p> <p>If the socket has not already been bound to a local address, <code>connect()</code> will bind it to an address which, unless the socket's address family is <code>AF_UNIX</code>, is an unused local address.</p> <p>If the initiating socket is not connection-mode, then <code>connect()</code> sets the socket's peer address, but no connection is made. For <code>SOCK_DGRAM</code> sockets, the peer address identifies where all datagrams are sent on subsequent <code>send(3XNET)</code> calls, and limits the remote sender for subsequent <code>recv(3XNET)</code> calls. If <i>address</i> is a null address for the protocol, the socket's peer address will be reset.</p> <p>If the initiating socket is connection-mode, then <code>connect()</code> attempts to establish a connection to the address specified by the <i>address</i> argument.</p> <p>If the connection cannot be established immediately and <code>O_NONBLOCK</code> is not set for the file descriptor for the socket, <code>connect()</code> will block for up to an unspecified timeout interval until the connection is established. If the timeout interval expires before the connection is established, <code>connect()</code> will fail and the connection attempt will be aborted. If <code>connect()</code> is interrupted by a signal that is caught while blocked waiting to establish a connection, <code>connect()</code> will fail and set <code>errno</code> to <code>EINTR</code>, but the connection request will not be aborted, and the connection will be established asynchronously.</p> <p>If the connection cannot be established immediately and <code>O_NONBLOCK</code> is set for the file descriptor for the socket, <code>connect()</code> will fail and set <code>errno</code> to <code>EINPROGRESS</code>, but the connection request will not be aborted, and the connection will be established asynchronously. Subsequent calls to <code>connect()</code> for the same socket, before the connection is established, will fail and set <code>errno</code> to <code>EALREADY</code>.</p> <p>When the connection has been established asynchronously, <code>select(3C)</code> and <code>poll(2)</code> will indicate that the file descriptor for the socket is ready for writing.</p>

## connect(3XNET)

The socket in use may require the process to have appropriate privileges to use the `connect ()` function.

**USAGE** If `connect ()` fails, the state of the socket is unspecified. Portable applications should close the file descriptor and create a new socket before attempting to reconnect.

**RETURN VALUES** Upon successful completion, `connect ()` returns 0. Otherwise, `-1` is returned and `errno` is set to indicate the error.

**ERRORS** The `connect ()` function will fail if:

<code>EADDRNOTAVAIL</code>	The specified address is not available from the local machine.
<code>EAFNOSUPPORT</code>	The specified address is not a valid address for the address family of the specified socket.
<code>EALREADY</code>	A connection request is already in progress for the specified socket.
<code>EBADF</code>	The <i>socket</i> argument is not a valid file descriptor.
<code>ECONNREFUSED</code>	The target address was not listening for connections or refused the connection request.
<code>EFAULT</code>	The address parameter can not be accessed.
<code>EINPROGRESS</code>	<code>O_NONBLOCK</code> is set for the file descriptor for the socket and the connection cannot be immediately established; the connection will be established asynchronously.
<code>EINTR</code>	The attempt to establish a connection was interrupted by delivery of a signal that was caught; the connection will be established asynchronously.
<code>EISCONN</code>	The specified socket is connection-mode and is already connected.
<code>ENETUNREACH</code>	No route to the network is present.
<code>ENOTSOCK</code>	The <i>socket</i> argument does not refer to a socket.
<code>EPROTOTYPE</code>	The specified address has a different type than the socket bound to the specified peer address.
<code>ETIMEDOUT</code>	The attempt to connect timed out before a connection was made.
If the address family of the socket is <code>AF_UNIX</code> , then <code>connect ()</code> will fail if:	
<code>EIO</code>	An I/O error occurred while reading from or writing to the file system.

connect(3XNET)

- ELOOP Too many symbolic links were encountered in translating the pathname in *address*.
- ENAMETOOLONG A component of a pathname exceeded `NAME_MAX` characters, or an entire pathname exceeded `PATH_MAX` characters.
- ENOENT A component of the pathname does not name an existing file or the pathname is an empty string.
- ENOTDIR A component of the path prefix of the pathname in *address* is not a directory.

The `connect()` function may fail if:

- EACCES Search permission is denied for a component of the path prefix; or write access to the named socket is denied.
- EADDRINUSE Attempt to establish a connection that uses addresses that are already in use.
- ECONNRESET Remote host reset the connection request.
- EHOSTUNREACH The destination host cannot be reached (probably because the host is down or a remote router cannot reach it).
- EINVAL The *address\_len* argument is not a valid length for the address family; or invalid address family in `sockaddr` structure.
- ENAMETOOLONG Pathname resolution of a symbolic link produced an intermediate result whose length exceeds `PATH_MAX`.
- ENETDOWN The local interface used to reach the destination is down.
- ENOBUFS No buffer space is available.
- ENOSR There were insufficient STREAMS resources available to complete the operation.
- EOPNOTSUPP The socket is listening and can not be connected.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO**

`close(2)`, `poll(2)`, `accept(3XNET)`, `bind(3XNET)`, `getsockname(3XNET)`, `select(3C)`, `send(3XNET)`, `shutdown(3XNET)`, `socket(3XNET)`, `attributes(5)`

## dial(3NSL)

<b>NAME</b>	dial – establish an outgoing terminal line connection										
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;dial.h&gt;  int dial(CALL call); void undial(int fd);</pre>										
<b>DESCRIPTION</b>	<p>dial() returns a file-descriptor for a terminal line open for read/write. The argument to dial() is a CALL structure (defined in the header &lt;dial.h&gt;).</p> <p>When finished with the terminal line, the calling program must invoke undial() to release the semaphore that has been set during the allocation of the terminal device.</p> <p>CALL is defined in the header &lt;dial.h&gt; and has the following members:</p> <pre>struct termio *attr; /* pointer to termio attribute struct */ int baud; /* transmission data rate */ int speed; /* 212A modem: low=300, high=1200 */ char *line; /* device name for out-going line */ char *telno; /* pointer to tel-no digits string */ int modem; /* specify modem control for direct lines */ char *device; /* unused */ int dev_len; /* unused */</pre> <p>The CALL element speed is intended only for use with an outgoing dialed call, in which case its value should be the desired transmission baud rate. The CALL element baud is no longer used.</p> <p>If the desired terminal line is a direct line, a string pointer to its device-name should be placed in the line element in the CALL structure. Legal values for such terminal device names are kept in the Devices file. In this case, the value of the baud element should be set to -1. This value will cause dial to determine the correct value from the &lt;Devices&gt; file.</p> <p>The telno element is for a pointer to a character string representing the telephone number to be dialed. Such numbers may consist only of these characters:</p> <table><tr><td>0-9</td><td>dial 0-9</td></tr><tr><td>*</td><td>dail *</td></tr><tr><td>#</td><td>dail #</td></tr><tr><td>=</td><td>wait for secondary dial tone</td></tr><tr><td>-</td><td>delay for approximately 4 seconds</td></tr></table> <p>The CALL element modem is used to specify modem control for direct lines. This element should be non-zero if modem control is required. The CALL element attr is a pointer to a termio structure, as defined in the header &lt;termio.h&gt;. A NULL value</p>	0-9	dial 0-9	*	dail *	#	dail #	=	wait for secondary dial tone	-	delay for approximately 4 seconds
0-9	dial 0-9										
*	dail *										
#	dail #										
=	wait for secondary dial tone										
-	delay for approximately 4 seconds										

for this pointer element may be passed to the `dial` function, but if such a structure is included, the elements specified in it will be set for the outgoing terminal line before the connection is established. This setting is often important for certain attributes such as parity and baud-rate.

The `CALL` elements `device` and `dev_len` are no longer used. They are retained in the `CALL` structure for compatibility reasons.

**RETURN VALUES**

On failure, a negative value indicating the reason for the failure will be returned. Mnemonics for these negative indices as listed here are defined in the header `<dial.h>`.

```
INTRPT  -1      /* interrupt occurred */
D_HUNG  -2      /* dialer hung (no return from write) */
NO_ANS  -3      /* no answer within 10 seconds */
ILL_BD  -4      /* illegal baud-rate */
A_PROB  -5      /* acu problem (open( ) failure) */
L_PROB  -6      /* line problem (open( ) failure) */
NO_LdV  -7      /* can't open Devices file */
DV_NT_A -8      /* requested device not available */
DV_NT_K -9      /* requested device not known */
NO_BD_A -10     /* no device available at requested baud */
NO_BD_K -11     /* no device known at requested baud */
DV_NT_E -12     /* requested speed does not match */
BAD_SYS -13     /* system not in Systems file*/
```

**FILES**

```
/etc/uucp/Devices
/etc/uucp/Systems
/var/spool/uucp/LCK..tty-device
```

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO**

`uucp(1C)`, `alarm(2)`, `read(2)`, `write(2)`, `attributes(5)`, `termio(7I)`

**NOTES**

Including the header `<dial.h>` automatically includes the header `<termio.h>`. An `alarm(2)` system call for 3600 seconds is made (and caught) within the `dial` module for the purpose of “touching” the `LCK.. .` file and constitutes the device allocation semaphore for the terminal device. Otherwise, `uucp(1C)` may simply delete the `LCK.. .` entry on its 90-minute clean-up rounds. The alarm may go off while the user program is in a `read(2)` or `write(2)` function, causing an apparent error return. If the user program expects to be around for an hour or more, error returns from `read( )`s should be checked for (`errno==EINTR`), and the `read( )` possibly reissued.

This interface is unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.

## doconfig(3NSL)

<b>NAME</b>	doconfig – execute a configuration script
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lnsl [ <i>library</i> ... ] # include &lt;sac.h&gt;  int doconfig(int <i>fildev</i>, char *<i>script</i>, long <i>rflag</i>);</pre>
<b>DESCRIPTION</b>	<p>doconfig() is a Service Access Facility library function that interprets the configuration scripts contained in the files &lt;/etc/saf/<i>pmtag</i>/<i>_config</i>&gt;, &lt;/etc/saf/<i>_sysconfig</i>&gt;, and &lt;/etc/saf/<i>pmtag</i>/<i>svctag</i>&gt;, where <i>pmtag</i> specifies the tag associated with the port monitor, and <i>svctag</i> specifies the service tag associated with a given service. See pmadm(1M) and sacadm(1M).</p> <p><i>script</i> is the name of the configuration script; <i>fildev</i> is a file descriptor that designates the stream to which stream manipulation operations are to be applied; <i>rflag</i> is a bitmask that indicates the mode in which <i>script</i> is to be interpreted. If <i>rflag</i> is zero, all commands in the configuration script are eligible to be interpreted. If <i>rflag</i> has the NOASSIGN bit set, the assign command is considered illegal and will generate an error return. If <i>rflag</i> has the NORUN bit set, the run and runwait commands are considered illegal and will generate error returns.</p> <p>The configuration language in which <i>script</i> is written consists of a sequence of commands, each of which is interpreted separately. The following reserved keywords are defined: assign, push, pop, runwait, and run. The comment character is #; when a # occurs on a line, everything from that point to the end of the line is ignored. Blank lines are not significant. No line in a command script may exceed 1024 characters.</p> <p><b>assign <i>variable=value</i></b> Used to define environment variables. <i>variable</i> is the name of the environment variable and <i>value</i> is the value to be assigned to it. The value assigned must be a string constant; no form of parameter substitution is available. <i>value</i> may be quoted. The quoting rules are those used by the shell for defining environment variables. assign will fail if space cannot be allocated for the new variable or if any part of the specification is invalid.</p> <p><b>push <i>module1</i> [, <i>module2</i>, <i>module3</i>, ...]</b> Used to push STREAMS modules onto the stream designated by <i>fildev</i>. <i>module1</i> is the name of the first module to be pushed, <i>module2</i> is the name of the second module to be pushed, etc. The command will fail if any of the named modules cannot be pushed. If a module cannot be pushed, the subsequent modules on the same command line will be ignored and modules that have already been pushed will be popped.</p> <p><b>pop [<i>module</i>]</b> Used to pop STREAMS modules off the designated stream. If pop is invoked with no arguments, the top module on the stream is popped. If an argument is given, modules will be popped one at a time until the named module is at the top of the stream. If the named module is not on the designated stream, the stream is left as it</p>



was and the command fails. If *module* is the special keyword ALL, then all modules on the stream will be popped. Note that only modules above the topmost driver are affected.

#### runwait command

The `runwait` command runs a command and waits for it to complete. `command` is the pathname of the command to be run. The command is run with `/usr/bin/sh -c` prepended to it; shell scripts may thus be executed from configuration scripts. The `runwait` command will fail if `command` cannot be found or cannot be executed, or if `command` exits with a non-zero status.

#### run command

The `run` command is identical to `runwait` except that it does not wait for `command` to complete. `command` is the pathname of the command to be run. `run` will not fail unless it is unable to create a child process to execute the command.

Although they are syntactically indistinguishable, some of the commands available to `run` and `runwait` are interpreter built-in commands. Interpreter built-ins are used when it is necessary to alter the state of a process within the context of that process. The `doconfig()` interpreter built-in commands are similar to the shell special commands and, like these, they do not spawn another process for execution. See `sh(1)`. The built-in commands are:

```
cd
ulimit
umask
```

#### RETURN VALUES

`doconfig()` returns 0 if the script was interpreted successfully. If a command in the script fails, the interpretation of the script ceases at that point and a positive number is returned; this number indicates which line in the script failed. If a system error occurs, a value of -1 is returned. When a script fails, the process whose environment was being established should *not* be started.

#### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

#### SEE ALSO

`sh(1)`, `pmadm(1M)`, `sacadm(1M)`, `attributes(5)`

#### NOTES

This interface is unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.

## endhostent(3XNET)

<b>NAME</b>	endhostent, gethostbyaddr, gethostbyname, gethostent, sethostent – network host database functions
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;netdb.h&gt; extern int h_errno;  void <b>endhostent</b>(void);  struct hostent *<b>gethostbyaddr</b>(const void *<i>addr</i>, size_t <i>len</i>, int     <i>type</i>);  struct hostent *<b>gethostbyname</b>(const char *<i>name</i>);  struct hostent *<b>gethostent</b>(void);  void <b>sethostent</b>(int <i>stayopen</i>);</pre>
<b>DESCRIPTION</b>	<p>The <code>gethostent()</code>, <code>gethostbyaddr()</code>, and <code>gethostbyname()</code> functions each return a pointer to a <code>hostent</code> structure, the members of which contain the fields of an entry in the network host database.</p> <p>The <code>gethostent()</code> function reads the next entry of the database, opening a connection to the database if necessary.</p> <p>The <code>gethostbyaddr()</code> function searches the database and finds an entry which matches the address family specified by the <code>type</code> argument and which matches the address pointed to by the <code>addr</code> argument, opening a connection to the database if necessary. The <code>addr</code> argument is a pointer to the binary-format (that is, not null-terminated) address in network byte order, whose length is specified by the <code>len</code> argument. The datatype of the address depends on the address family. For an address of type <code>AF_INET</code>, this is an <code>in_addr</code> structure, defined in <code>&lt;netinet/in.h&gt;</code>. For an address of type <code>AF_INET6</code>, there is an <code>in6_addr</code> structure defined in <code>&lt;netinet/in.h&gt;</code>.</p> <p>The <code>gethostbyname()</code> function searches the database and finds an entry which matches the host name specified by the <code>name</code> argument, opening a connection to the database if necessary. If <code>name</code> is an alias for a valid host name, the function returns information about the host name to which the alias refers, and <code>name</code> is included in the list of aliases returned.</p> <p>The <code>sethostent()</code> function opens a connection to the network host database, and sets the position of the next entry to the first entry. If the <code>stayopen</code> argument is non-zero, the connection to the host database will not be closed after each call to <code>gethostent()</code> (either directly, or indirectly through one of the other <code>gethost*()</code> functions).</p> <p>The <code>endhostent()</code> function closes the connection to the database.</p>
<b>USAGE</b>	The <code>gethostent()</code> , <code>gethostbyaddr()</code> , and <code>gethostbyname()</code> functions may return pointers to static data, which may be overwritten by subsequent calls to any of these functions.

These functions are generally used with the Internet address family.

**RETURN VALUES**

On successful completion, `gethostbyaddr()`, `gethostbyname()` and `gethostent()` return a pointer to a `hostent` structure if the requested entry was found, and a null pointer if the end of the database was reached or the requested entry was not found. Otherwise, a null pointer is returned.

On unsuccessful completion, `gethostbyaddr()` and `gethostbyname()` functions set `h_errno` to indicate the error.

**ERRORS**

No errors are defined for `endhostent()`, `gethostent()` and `sethostent()`.

The `gethostbyaddr()` and `gethostbyname()` functions will fail in the following cases, setting `h_errno` to the value shown in the list below. Any changes to `errno` are unspecified.

HOST_NOT_FOUND	No such host is known.
NO_DATA	The server recognised the request and the name but no address is available. Another type of request to the name server for the domain might return an answer.
NO_RECOVERY	An unexpected server failure occurred which can not be recovered.
TRY_AGAIN	A temporary and possibly transient error occurred, such as a failure of a server to respond.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO**

`endservent(3XNET)`, `htonl(3XNET)`, `inet_addr(3XNET)`, `attributes(5)`

## endnetent(3XNET)

<b>NAME</b>	endnetent, getnetbyaddr, getnetbyname, getnetent, setnetent – network database functions
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;netdb.h&gt;  void <b>endnetent</b>(void); struct netent *getnetbyaddr(in_addr_t <i>net</i>, int <i>type</i>); struct netent *<b>getnetbyname</b>(const char *<i>name</i>); struct netent *<b>getnetent</b>(void); void <b>setnetent</b>(int <i>stayopen</i>);</pre>
<b>DESCRIPTION</b>	<p>The <code>getnetbyaddr()</code>, <code>getnetbyname()</code> and <code>getnetent()</code>, functions each return a pointer to a <code>netent</code> structure, the members of which contain the fields of an entry in the network database.</p> <p>The <code>getnetent()</code> function reads the next entry of the database, opening a connection to the database if necessary.</p> <p>The <code>getnetbyaddr()</code> function searches the database from the beginning, and finds the first entry for which the address family specified by <code>type</code> matches the <code>n_addrtype</code> member and the network number <code>net</code> matches the <code>n_net</code> member, opening a connection to the database if necessary. The <code>net</code> argument is the network number in host byte order.</p> <p>The <code>getnetbyname()</code> function searches the database from the beginning and finds the first entry for which the network name specified by <code>name</code> matches the <code>n_name</code> member, opening a connection to the database if necessary.</p> <p>The <code>setnetent()</code> function opens and rewinds the database. If the <code>stayopen</code> argument is non-zero, the connection to the net database will not be closed after each call to <code>getnetent()</code> (either directly, or indirectly through one of the other <code>getnet*( )</code> functions).</p> <p>The <code>endnetent()</code> function closes the database.</p>
<b>USAGE</b>	<p>The <code>getnetbyaddr()</code>, <code>getnetbyname()</code> and <code>getnetent()</code>, functions may return pointers to static data, which may be overwritten by subsequent calls to any of these functions.</p> <p>These functions are generally used with the Internet address family.</p>
<b>RETURN VALUES</b>	On successful completion, <code>getnetbyaddr()</code> , <code>getnetbyname()</code> and <code>getnetent()</code> , return a pointer to a <code>netent</code> structure if the requested entry was found, and a null pointer if the end of the database was reached or the requested entry was not found. Otherwise, a null pointer is returned.
<b>ERRORS</b>	No errors are defined.

endnetent(3XNET)

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO** `attributes(5)`

## endprotoent(3XNET)

<b>NAME</b>	endprotoent, getprotobynumber, getprotobynname, getprotoent, setprotoent – network protocol database functions
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;netdb.h&gt;  void <b>endprotoent</b> (void) ;  struct protoent *<b>getprotobynname</b> (const char *<i>name</i>) ;  struct protoent *<b>getprotobynumber</b> (int <i>proto</i>) ;  struct protoent *<b>getprotoent</b> (void) ;  void <b>setprotoent</b> (int <i>stayopen</i>) ;</pre>
<b>DESCRIPTION</b>	<p>The <code>getprotobynname()</code>, <code>getprotobynumber()</code> and <code>getprotoent()</code>, functions each return a pointer to a <code>protoent</code> structure, the members of which contain the fields of an entry in the network protocol database.</p> <p>The <code>getprotoent()</code> function reads the next entry of the database, opening a connection to the database if necessary.</p> <p>The <code>getprotobynname()</code> function searches the database from the beginning and finds the first entry for which the protocol name specified by <i>name</i> matches the <code>p_name</code> member, opening a connection to the database if necessary.</p> <p>The <code>getprotobynumber()</code> function searches the database from the beginning and finds the first entry for which the protocol number specified by <i>number</i> matches the <code>p_proto</code> member, opening a connection to the database if necessary.</p> <p>The <code>setprotoent()</code> function opens a connection to the database, and sets the next entry to the first entry. If the <i>stayopen</i> argument is non-zero, the connection to the network protocol database will not be closed after each call to <code>getprotoent()</code> (either directly, or indirectly through one of the other <code>getproto*()</code> functions).</p> <p>The <code>endprotoent()</code> function closes the connection to the database.</p>
<b>USAGE</b>	<p>The <code>getprotobynname()</code>, <code>getprotobynumber()</code> and <code>getprotoent()</code> functions may return pointers to static data, which may be overwritten by subsequent calls to any of these functions.</p> <p>These functions are generally used with the Internet address family.</p>
<b>RETURN VALUES</b>	On successful completion, <code>getprotobynname()</code> , <code>getprotobynumber()</code> and <code>getprotoent()</code> functions return a pointer to a <code>protoent</code> structure if the requested entry was found, and a null pointer if the end of the database was reached or the requested entry was not found. Otherwise, a null pointer is returned.
<b>ERRORS</b>	No errors are defined.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

endprotoent(3XNET)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO** attributes(5)

## endservent(3XNET)

<b>NAME</b>	endservent, getservbyport, getservbyname, getservent, setservent – network services database functions
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;netdb.h&gt;  void endservent(void);  struct servent *getservbyname(const char *name, const char *proto); struct servent *getservbyport(int port, const char *proto); struct servent *getservent(void); void setservent(int stayopen);</pre>
<b>DESCRIPTION</b>	<p>The <code>getservbyname()</code>, <code>getservbyport()</code> and <code>getservent()</code> functions each return a pointer to a <code>servent</code> structure, the members of which contain the fields of an entry in the network services database.</p> <p>The <code>getservent()</code> function reads the next entry of the database, opening a connection to the database if necessary.</p> <p>The <code>getservbyname()</code> function searches the database from the beginning and finds the first entry for which the service name specified by <i>name</i> matches the <code>s_name</code> member and the protocol name specified by <i>proto</i> matches the <code>s_proto</code> member, opening a connection to the database if necessary. If <i>proto</i> is a null pointer, any value of the <code>s_proto</code> member will be matched.</p> <p>The <code>getservbyport()</code> function searches the database from the beginning and finds the first entry for which the port specified by <i>port</i> matches the <code>s_port</code> member and the protocol name specified by <i>proto</i> matches the <code>s_proto</code> member, opening a connection to the database if necessary. If <i>proto</i> is a null pointer, any value of the <code>s_proto</code> member will be matched. The <i>port</i> argument must be in network byte order.</p> <p>The <code>setservent()</code> function opens a connection to the database, and sets the next entry to the first entry. If the <i>stayopen</i> argument is non-zero, the net database will not be closed after each call to the <code>getservent()</code> function (either directly, or indirectly through one of the other <code>getserv*()</code> functions).</p> <p>The <code>endservent()</code> function closes the database.</p>
<b>USAGE</b>	<p>The <i>port</i> argument of <code>getservbyport()</code> need not be compatible with the port values of all address families.</p> <p>The <code>getservent()</code>, <code>getservbyname()</code> and <code>getservbyport()</code> functions may return pointers to static data, which may be overwritten by subsequent calls to any of these functions.</p> <p>These functions are generally used with the Internet address family.</p>



endservent(3XNET)

**RETURN VALUES** On successful completion, `getservbyname()`, `getservbyport()` and `getservent()` return a pointer to a `servent` structure if the requested entry was found, and a null pointer if the end of the database was reached or the requested entry was not found. Otherwise, a null pointer is returned.

**ERRORS** No errors are defined.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO** `endhostent(3XNET)`, `endprotoent(3XNET)`, `htonl(3XNET)`, `inet_addr(3XNET)`, `attributes(5)`

## ethers(3SOCKET)

<b>NAME</b>	ethers, ether_ntoa, ether_aton, ether_ntohost, ether_hostton, ether_line – Ethernet address mapping operations
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt; #include &lt;net/if.h&gt; #include &lt;netinet/in.h&gt; #include &lt;netinet/if_ether.h&gt;  char *ether_ntoa(struct ether_addr *e); struct ether_addr *ether_aton(char *s); int ether_ntohost(char *hostname, struct ether_addr *e); int ether_hostton(char *hostname, struct ether_addr *e); int ether_line(char *l, struct ether_addr *e, char *hostname);</pre>
<b>DESCRIPTION</b>	<p>These routines are useful for mapping 48 bit Ethernet numbers to their ASCII representations or their corresponding host names, and vice versa.</p> <p>The function <code>ether_ntoa()</code> converts a 48 bit Ethernet number pointed to by <code>e</code> to its standard ASCII representation; it returns a pointer to the ASCII string. The representation is of the form <code>x:x:x:x:x:x</code> where <code>x</code> is a hexadecimal number between 0 and <code>ff</code>. The function <code>ether_aton()</code> converts an ASCII string in the standard representation back to a 48 bit Ethernet number; the function returns <code>NULL</code> if the string cannot be scanned successfully.</p> <p>The function <code>ether_ntohost()</code> maps an Ethernet number (pointed to by <code>e</code>) to its associated hostname. The string pointed to by <code>hostname</code> must be long enough to hold the hostname and a <code>NULL</code> character. The function returns zero upon success and non-zero upon failure. Inversely, the function <code>ether_hostton()</code> maps a hostname string to its corresponding Ethernet number; the function modifies the Ethernet number pointed to by <code>e</code>. The function also returns zero upon success and non-zero upon failure. In order to do the mapping, both these functions may lookup one or more of the following sources: the <code>ethers</code> file, the NIS maps “<code>ethers.byname</code>” and “<code>ethers.byaddr</code>” and the NIS+ table “<code>ethers</code>”. The sources and their lookup order are specified in the <code>/etc/nsswitch.conf</code> file (see <code>nsswitch.conf(4)</code> for details).</p> <p>The function <code>ether_line()</code> scans a line (pointed to by <code>l</code>) and sets the hostname and the Ethernet number (pointed to by <code>e</code>). The string pointed to by <code>hostname</code> must be long enough to hold the hostname and a <code>NULL</code> character. The function returns zero upon success and non-zero upon failure. The format of the scanned line is described by <code>ethers(4)</code>.</p>

ethers(3SOCKET)

**FILES** /etc/ethers

/etc/nsswitch.conf

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** ethers(4), nsswitch.conf(4), attributes(5)

**BUGS** Programs that call `ether_hostton()` or `ether_ntohost()` routines cannot be linked statically since the implementation of these routines requires dynamic linker functionality to access shared objects at run time.

## fn\_attr\_bind(3XFN)

**NAME** | `fn_attr_bind` – bind a reference to a name and associate attributes with named object

**SYNOPSIS** | `#include <xfn/xfn.h>`

```
int fn_attr_bind(FN_ctx_t *ctx, const FN_composite_name_t *name,
                  const FN_ref_t *ref, const FN_attrset_t *attrs, unsigned int
                  exclusive, FN_status_t *status);
```

**DESCRIPTION** | This operation binds the supplied reference *ref* to the supplied composite name *name* relative to *ctx*, and associates the attributes specified in *attrs* with the named object. The binding is made in the target context, that is, that context named by all but the terminal atomic part of *name*. The operation binds the terminal atomic name to the supplied reference in the target context. The target context must already exist.

The value of *exclusive* determines what happens if the terminal atomic part of the name is already bound in the target context. If *exclusive* is nonzero and *name* is already bound, the operation fails. If *exclusive* is 0, the new binding replaces any existing binding, and, if *attrs* is not NULL, *attrs* replaces any existing attributes associated with the named object. If *attrs* is NULL and *exclusive* is 0, any existing attributes associated with the named object are left unchanged.

**RETURN VALUES** | `fn_attr_bind()` returns 1 upon success, 0 upon failure.

**ERRORS** | `fn_attr_bind()` sets *status* as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`. Of special relevance for this operation is the following status code:

`FN_E_NAME_IN_USE`            The supplied name is already in use.

**USAGE** | The value of *ref* cannot be NULL. If the intent is to reserve a name using `fn_attr_bind()`, a reference containing no address should be supplied. This reference may be name service-specific or it may be the conventional NULL reference.

If multiple sources are updating a reference or attributes associated with a named object, they must synchronize amongst each other when adding, modifying, or removing from the address list of a bound reference, or manipulating attributes associated with the named object.

**ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** | `FN_composite_name_t(3XFN)`, `FN_ctx_t(3XFN)`, `FN_ref_t(3XFN)`, `FN_status_t(3XFN)`, `fn_ctx_bind(3XFN)`, `fn_ctx_lookup(3XFN)`, `fn_ctx_unbind(3XFN)`, `xfn_attributes(3XFN)`, `xfn_status_codes(3XFN)`, `attributes(5)`

<b>NAME</b>	fn_attr_create_subcontext – create a subcontext in a context and associate attributes with newly created context				
<b>SYNOPSIS</b>	<pre>#include &lt;xfn/xfn.h&gt;  FN_ref_t *fn_attr_create_subcontext(FN_ctx_t *ctx, const     FN_composite_name_t *name, const FN_attrset_t *attrs,     FN_status_t *status);</pre>				
<b>DESCRIPTION</b>	<p>This operation creates a new XFN context of the same type as the target context, that is, that context named by all but the terminal atomic component of <i>name</i>, and binds it to the supplied composite name. In addition, attributes given in <i>attrs</i> are associated with the newly created context.</p> <p>The target context must already exist. The new context is created and bound in the target context using the terminal atomic name in <i>name</i>. The operation returns a reference to the newly created context.</p>				
<b>RETURN VALUES</b>	fn_attr_create_subcontext() returns a reference to the newly created context; if the operation fails, it returns a NULL pointer.				
<b>ERRORS</b>	<p>fn_attr_create_subcontext() sets <i>status</i> as described in FN_status_t(3XFN) and xfn_status_codes(3XFN). Of special relevance for this operation is the following status code:</p> <p>FN_E_NAME_IN_USE            The terminal atomic name already exists in the target context.</p>				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	FN_composite_name_t(3XFN), FN_ctx_t(3XFN), FN_ref_t(3XFN), FN_status_t(3XFN), fn_attr_bind(3XFN), fn_ctx_bind(3XFN), fn_ctx_create_subcontext(3XFN), fn_ctx_destroy_subcontext(3XFN), fn_ctx_lookup(3XFN), xfn_attributes(3XFN), xfn_status_codes(3XFN), attributes(5)				

## fn\_attr\_ext\_search(3XFN)

<b>NAME</b>	fn_attr_ext_search, FN_ext_searchlist_t, fn_ext_searchlist_next, fn_ext_searchlist_destroy – search for names in the specified context(s) whose attributes satisfy the filter
<b>SYNOPSIS</b>	<pre>#include &lt;xfn/xfn.h&gt;  FN_ext_searchlist_t *fn_attr_ext_search(FN_ctx_t *ctx, const     FN_composite_name_t *name, const FN_search_control_t *control,     const FN_search_filter_t *filter, FN_status_t *status);  FN_composite_name_t *fn_ext_searchlist_next(FN_ext_searchlist_t     *esl, FN_ref_t **returned_ref, FN_attrset_t **returned_attrs,     FN_status_t *status);  void fn_ext_searchlist_destroy(FN_ext_searchlist_t *esl);</pre>
<b>DESCRIPTION</b>	<p>This set of operations is used to list names of objects whose attributes satisfy the filter expression. The references to which these names are bound and specified attributes and their values may also be returned.</p> <p><i>control</i> encapsulates the option settings for the search. These options are:</p> <ul style="list-style-type: none"><li>■ the scope of the search</li><li>■ whether XFN links are followed</li><li>■ a limit on the number of names returned</li><li>■ whether references and specific attributes associated with the named objects that satisfy the filter are returned</li></ul> <p>The scope of the search is one of:</p> <ul style="list-style-type: none"><li>■ the object named <i>name</i> relative to the context <i>ctx</i></li><li>■ the context named <i>name</i> relative to the context <i>ctx</i></li><li>■ the context named <i>name</i> relative to the context <i>ctx</i>, and its subcontexts</li></ul> <p>or</p> <ul style="list-style-type: none"><li>■ the context named <i>name</i> relative to the context <i>ctx</i>, and a context implementation-defined set of subcontexts</li></ul> <p>If the value of <i>control</i> is 0, default control option settings are used. The default settings are:</p> <ul style="list-style-type: none"><li>■ scope is search named context</li><li>■ links are not followed</li><li>■ all names of objects that satisfy the filter are returned</li><li>■ references and attributes are not returned</li></ul> <p>The <code>FN_search_control_t</code> type is described in <code>FN_search_control_t(3XFN)</code>.</p>

fn\_attr\_ext\_search(3XFN)

The filter expression *filter* in `fn_attr_ext_search()` is evaluated against the attributes of the objects bound in the scope of the search. The filter evaluates to either TRUE or FALSE. The names and, optionally, the references and attributes of objects whose attributes satisfy the filter are enumerated. If the value of *filter* is 0, all names within the search scope are enumerated. The `FN_search_filter_t` type is described in `FN_search_filter_t(3XFN)`.

The call to `fn_attr_ext_search()` initiates the search process. It returns a handle to an `FN_ext_searchlist_t` object that is used to enumerate the names of the objects that satisfy the filter.

The operation `fn_ext_searchlist_next()` returns the next name in the enumeration identified by *esl*; it also updates *esl* to indicate the state of the enumeration. If the reference to which the name is bound was requested, it is returned in *returned\_ref*. Requested attributes associated with the name are returned in *returned\_attrs*; each attribute consists of an attribute identifier, syntax, and value(s). Successive calls to `fn_ext_searchlist_next()` using *esl* return successive names and, optionally, their references and attributes, in the enumeration; these calls further update the state of the enumeration.

The names that are returned are composite names, to be resolved relative to the starting context for the search. This starting context is the context named *name* relative to *ctx* unless the scope of the search is only the named object. If the scope of the search is only the named object, the terminal atomic name in *name* is returned.

`fn_ext_searchlist_destroy()` releases resources used during the enumeration. This may be invoked at any time to terminate the enumeration.

## RETURN VALUES

`fn_attr_ext_search()` returns a pointer to an `FN_ext_searchlist_t` object if the search is successfully initiated; it returns a NULL pointer if the search cannot be initiated or if no named object with attributes whose values satisfy the filter expression is found.

`fn_ext_searchlist_next()` returns a pointer to an `FN_composite_name_t` object (see `FN_composite_name_t(3XFN)`) that is the next name in the enumeration; it returns a NULL pointer if no more names can be returned. If *returned\_attrs* is a NULL pointer, no attributes are returned; otherwise, *returned\_attrs* contains the attributes associated with the named object, as specified in the control parameter to `fn_attr_ext_search()`. If *returned\_ref* is a NULL pointer, no reference is returned; otherwise, if *control* specified the return of the reference of the named object, that reference is returned in *returned\_ref*.

In the case of a failure, these operations return in the *status* argument a code indicating the nature of the failure.

## ERRORS

If successful, `fn_attr_ext_search()` returns a pointer to an `FN_ext_searchlist_t` object and sets *status* to `FN_SUCCESS`.

`fn_attr_ext_search()` returns a NULL pointer when no more names can be returned. *status* is set in the following way:

## fn\_attr\_ext\_search(3XFN)

FN_SUCCESS	A named object could not be found whose attributes satisfied the filter expression.
FN_E_NOT_A_CONTEXT	The object named for the start of the search was not a context and the search scope was the given context or the given context and its subcontexts.
FN_E_SEARCH_INVALID_FILTER	The filter could not be evaluated TRUE or FALSE, or there was some other problem with the filter.
FN_E_SEARCH_INVALID_OPTION	A supplied search control option could not be supported.
FN_E_SEARCH_INVALID_OP	An operator in the filter expression is not supported or, if the operator is an extended operator, the number of types of arguments supplied does not match the signature of the operation.
FN_E_ATTR_NO_PERMISSION	The caller did not have permission to read one or more of the attributes specified in the filter.
FN_E_INVALID_ATTR_VALUE	A value type in the filter did not match the syntax of the attribute against which it was being evaluated.

Other status codes are possible as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`.

Each successful call to `fn_ext_searchlist_next()` returns a name and, optionally, its reference in `returned_ref` and requested attributes in `returned_attrs`. `status` is set in the following way:

FN_SUCCESS	All requested attributes were returned successfully with the name.
FN_E_ATTR_NO_PERMISSION	The caller did not have permission to read one or more of the requested attributes.
FN_E_INVALID_ATTR_IDENTIFIER	A requested attribute identifier was not in a format acceptable to the naming system, or its contents were not valid for the format specified.
FN_E_NO_SUCH_ATTRIBUTE	The named object did not have one of the requested attributes.
FN_E_INSUFFICIENT_RESOURCES	Insufficient resources are available to return all the requested attributes and their values.



fn\_attr\_ext\_search(3XFN)

FN_E_ATTR_NO_PERMISSION	
FN_E_INVALID_ATTR_IDENTIFIER	
FN_E_NO_SUCH_ATTRIBUTE	
FN_E_INSUFFICIENT_RESOURCES	These indicate that some of the requested attributes may have been returned in <i>returned_attrs</i> but one or more of them could not be returned. Use <code>fn_attr_get(3XFN)</code> or <code>fn_attr_multi_get(3XFN)</code> to discover why these attributes could not be returned.

If `fn_ext_searchlist_next()` returns a name, it can be called again to get the next name in the enumeration.

`fn_ext_searchlist_next()` returns a NULL pointer if no more names can be returned. *status* is set in the following way:

FN_SUCCESS	The search has completed successfully.
FN_E_PARTIAL_RESULT	The enumeration is not yet complete but cannot be continued.
FN_E_ATTR_NO_PERMISSION	The caller did not have permission to read one or more of the attributes specified in the filter.
FN_E_INVALID_ENUM_HANDLE	The supplied enumeration handle was not valid. Possible reasons could be that the handle was from another enumeration, or the context being enumerated no longer accepts the handle (due to such events as handle expiration or updates to the context).

Other status codes are possible as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`.

**USAGE** The search performed by `fn_attr_ext_search()` is not ordered in any way, including the traversal of subcontexts. The names enumerated using `fn_ext_searchlist_next()` are not ordered in any way. Furthermore, there is no guarantee that any two series of enumerations with the same arguments to `fn_attr_ext_search()` will return the names in the same order.

XFN links encountered during the resolution of *name* are followed, regardless of the follow links control setting, and the search starts at the final named object or context.

If *control* specifies that the search should follow links, XFN link names encountered during the search are followed and the terminal named object is searched. If the terminal named object is bound to a context and the scope of the search includes subcontexts, that context and its subcontexts are also searched. For example, if *aname*

## fn\_attr\_ext\_search(3XFN)

is bound to an XFN link, *lname*, in a context within the scope of the search, and *aname* is returned by `fn_ext_searchlist_next()`, this means that the object identified by *lname* satisfied the filter expression. *aname* is returned instead of *lname* because *aname* can always be named relative to the starting context for the search.

If *control* specifies that the search should not follow links, the attributes associated with the names of XFN links are searched. For example, if *aname* is bound to an XFN link, *lname*, in a context within the scope of the search, and *aname* is returned by `fn_ext_searchlist_next()`, this means that the object identified by *aname* satisfied the filter expression.

When following XFN links, `fn_attr_ext_search()` may search contexts outside of *scope*. In addition, if the link name's terminal atomic name is bound in a context within *scope*, the operation may return the same object more than once.

XFN does not specify how *control* affects the following of native naming system links during the search.

### EXAMPLES

**EXAMPLE 1** A sample program of displaying how the `fn_attr_ext_search()` operation may be used.

The following code fragment illustrates how the `fn_attr_ext_search()` operation may be used. The code consists of three parts: preparing the arguments for the search, performing the search, and cleaning up.

The first part involves getting the name of the context to start the search and constructing the search filter that named objects in the context must satisfy. This is done in the declarations part of the code and by the routine `get_search_query`. See `FN_search_filter_t(3XFN)` for the description of *sfilter* and the filter creation operation.

The next part involves doing the search and enumerating the results of the search. This is done by first getting a context handle to the Initial Context, and then passing that handle along with the name of the target context and search filter to `fn_attr_ext_search()`. This particular call to `fn_attr_ext_search()` uses the default search control options (by passing in 0 as the *control* argument). This means that the search will be performed in the context named by *target\_name* and that no reference or attributes will be returned. In addition, any XFN links encountered will not be followed and all named objects that satisfy the search filter will be returned (that is, no limit). If successful, `fn_attr_ext_search()` returns *esl*, a handle for enumerating the results of the search. The results of the search are enumerated using calls to `fn_ext_searchlist_next()`, which returns the name of the object. (The arguments *returned\_ref* and *returned\_attrs* to `fn_ext_searchlist_next()` are 0 because the default search control used in `fn_attr_ext_search()` did not request them to be returned.)

**EXAMPLE 1** A sample program of displaying how the `fn_attr_ext_search()` operation may be used. (Continued)

The last part of the code involves cleaning up the resources used during the search and enumeration. The call to `fn_ext_searchlist_destroy()` releases resources reserved for this enumeration. The other calls release the context handle, name, filter, and status objects created earlier.

```

/* Declarations */
FN_ctx_t *ctx;
FN_ext_searchlist_t *esl;
FN_composite_name_t *name;
FN_status_t *status = fn_status_create();
FN_composite_name_t *target_name = get_name_from_user_input();
FN_search_filter_t *sfilter = get_search_query();
/* Get context handle to Initial Context */
ctx = fn_ctx_handle_from_initial(status);
/* error checking on 'status' */
/* Initiate search */
if ((esl=fn_attr_ext_search(ctx, target_name,
/* default controls */ 0, sfilter, status)) == 0) {
/* report 'status', cleanup, and exit */
}
/* Enumerate names requested */
while (name=fn_ext_searchlist_next(esl, 0, 0, status)) {
/* do something with 'name' */
fn_composite_destroy(name);
}
/* check 'status' for reason for end of enumeration */
/* Clean up */
fn_ext_searchlist_destroy(esl);
fn_search_filter_destroy(sfilter);
fn_ctx_handle_destroy(ctx);
fn_composite_name_destroy(target_name);
fn_status_destroy(status);
/*
* Procedure for constructing the filter object for search:
* "age" attribute is greater than or equal to 17 AND
* less than or equal to 25
* AND the "student" attribute is present.
*/
FN_search_filter_t *
get_search_query()
{
extern FN_attribute_t *attr_age;
extern FN_attribute_t *attr_student;
FN_search_filter_t *sfilter;
unsigned int filter_status;
sfilter = fn_search_filter_create(
&filter_status,
"(%a >= 17) and (%a <= 25) and %a",
attr_age, attr_age, attr_student);
/* error checking on 'filter_status' */
return (sfilter);
}

```

fn\_attr\_ext\_search(3XFN)

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_attrset\_t(3XFN), FN\_composite\_name\_t(3XFN), FN\_ctx\_t(3XFN), FN\_ref\_t(3XFN), FN\_search\_control\_t(3XFN), FN\_search\_filter\_t(3XFN), FN\_status\_t(3XFN), fn\_attr\_get(3XFN), fn\_attr\_multi\_get(3XFN), xfn\_status\_codes(3XFN), attributes(5)

- NAME** | fn\_attr\_get – return specified attribute associated with name
- SYNOPSIS** | 

```
cc [ flag ... ] file ... -lxfn [ library ... ]
#include <xfn/xfn.h>

FN_attribute_t *fn_attr_get(FN_ctx_t *ctx, const
    FN_composite_name_t *name, const FN_identifier_t *attribute_id,
    unsigned int follow_link, FN_status_t *status);
```
- DESCRIPTION** | This operation returns the identifier, syntax and values of a specified attribute for the object named *name* relative to *ctx*. If *name* is empty, the attribute associated with *ctx* is returned.
- The value of *follow\_link* determines what happens when the terminal atomic part of *name* is bound to an XFN link. If *follow\_link* is non-zero, such a link is followed, and the values of the attribute associated with the final named object are returned; if *follow\_link* is zero, such a link is not followed. Any XFN links encountered before the terminal atomic name are always followed.
- RETURN VALUES** | fn\_attr\_get returns a pointer to an FN\_attribute\_t object if the operation succeeds; it returns a NULL pointer (0) if the operation fails.
- ERRORS** | fn\_attr\_get() sets *status* as described in FN\_status\_t(3XFN) and xfn\_status\_codes(3XFN).
- USAGE** | fn\_attr\_get\_values() and its related operations are used for getting individual values of an attribute. They should be used if the combined size of all the values are expected to be too large to be returned in a single invocation of fn\_attr\_get().
- ATTRIBUTES** | See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

- SEE ALSO** | FN\_attribute\_t(3XFN), FN\_composite\_name\_t(3XFN), FN\_ctx\_t(3XFN), FN\_identifier\_t(3XFN), FN\_status\_t(3XFN), fn\_attr\_get\_values(3XFN), xfn(3XFN), xfn\_attributes(3XFN), xfn\_status\_codes(3XFN), attributes(5)
- NOTES** | The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_attr\_get\_ids(3XFN)

<b>NAME</b>	fn_attr_get_ids – get a list of the identifiers of all attributes associated with named object				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_attrset_t *fn_attr_get_ids(FN_ctx_t *ctx, const     FN_composite_name_t *name, unsigned int follow_link, FN_status_t     *status);</pre>				
<b>DESCRIPTION</b>	<p>This operation returns a list of the attribute identifiers of all attributes associated with the object named by <i>name</i> relative to the context <i>ctx</i>. If <i>name</i> is empty, the attribute identifiers associated with <i>ctx</i> are returned.</p> <p>The value of <i>follow_link</i> determines what happens when the terminal atomic part of <i>name</i> is bound to an XFN link. If <i>follow_link</i> is non-zero, such a link is followed, and the values of the attribute associated with the final named object are returned; if <i>follow_link</i> is zero, such a link is not followed. Any XFN links encountered before the terminal atomic name are always followed.</p>				
<b>RETURN VALUES</b>	This operation returns a pointer to an object of type <code>FN_attrset_t</code> ; if the operation fails, a NULL pointer (0) is returned.				
<b>ERRORS</b>	This operation sets <i>status</i> as described in <code>FN_status_t(3XFN)</code> and <code>xfn_status_codes(3XFN)</code> .				
<b>USAGE</b>	The attributes in the returned set do not contain the syntax or values of the attributes, only their identifiers.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>FN_attribute_t(3XFN)</code> , <code>FN_attrset_t(3XFN)</code> , <code>FN_composite_name_t(3XFN)</code> , <code>FN_ctx_t(3XFN)</code> , <code>FN_status_t(3XFN)</code> , <code>fn_attr_get(3XFN)</code> , <code>fn_attr_multi_get(3XFN)</code> , <code>xfn(3XFN)</code> , <code>xfn_attributes(3XFN)</code> , <code>xfn_status_codes(3XFN)</code> , <code>attributes(5)</code>				
<b>NOTES</b>	The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.				

fn\_attr\_get\_values(3XFN)

<b>NAME</b>	fn_attr_get_values, FN_valuelist_t, fn_valuelist_next, fn_valuelist_destroy – return values of an attribute
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_valuelist_t *fn_attr_get_values(FN_ctx_t *ctx, const     FN_composite_name_t *name, const FN_identifier_t *attribute_id,     unsigned int follow_link, FN_status_t *status);  FN_attrvalue_t *fn_valuelist_next(FN_valuelist_t *vl,     FN_identifier_t **attr_syntax, FN_status_t *status);  void fn_valuelist_destroy(FN_valuelist_t *vl, FN_status_t *status);</pre>
<b>DESCRIPTION</b>	<p>This set of operations is used to obtain the values of a single attribute, identified by <i>attribute_id</i>, associated with the object named <i>name</i>, resolved in the context <i>ctx</i>. If <i>name</i> is empty, the attribute values associated with <i>ctx</i> are obtained.</p> <p>The value of <i>follow_link</i> determines what happens when the terminal atomic part of <i>name</i> is bound to an XFN link. If <i>follow_link</i> is non-zero, such a link is followed, and the values of the attribute associated with the final named object are returned; if <i>follow_link</i> is zero, such a link is not followed. Any XFN links encountered before the terminal atomic name are always followed.</p> <p>The operation <code>fn_attr_get_values()</code> initiates the enumeration process. It returns a handle to an <code>FN_valuelist_t</code> object that can be used to enumerate the values of the specified attribute.</p> <p>The operation <code>fn_valuelist_next()</code> returns a new <code>FN_attrvalue_t</code> object containing the next value in the attribute and may be called multiple times until all values are retrieved. The syntax of the attribute is returned in <i>attr_syntax</i>.</p> <p>The operation <code>fn_valuelist_destroy()</code> is used to release the resources used during the enumeration. This may be invoked before the enumeration has completed to terminate the enumeration.</p> <p>These operations work in a fashion similar to the <code>fn_ctx_list_names()</code> operations.</p>
<b>RETURN VALUES</b>	<p><code>fn_attr_get_values()</code> returns a pointer to an <code>FN_valuelist_t</code> object if the enumeration process is successfully initiated; it returns a NULL pointer if the process failed.</p> <p><code>fn_valuelist_next()</code> returns a NULL pointer if no more attribute values can be returned.</p> <p>In the case of a failure, these operations set <i>status</i> to indicate the nature of the failure.</p>
<b>ERRORS</b>	Each successful call to <code>fn_valuelist_next()</code> returns an attribute value. <i>status</i> is set to <code>FN_SUCCESS</code> .

## fn\_attr\_get\_values(3XFN)

When `fn_valuelist_next()` returns a NULL pointer, it indicates that no more values can be returned. *status* is set in the following way:

<code>FN_SUCCESS</code>	The enumeration has completed successfully.
<code>FN_E_INVALID_ENUM_HANDLE</code>	The given enumeration handle is not valid. Possible reasons could be that the handle was from another enumeration, or the context being enumerated no longer accepts the handle (due to such events as handle expiration or updates to the context).
<code>FN_E_PARTIAL_RESULT</code>	The enumeration is not yet complete but cannot be continued.

In addition to these status codes, other status codes are also possible in calls to these operations. In such cases, *status* is set as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`.

**USAGE** This interface should be used instead of `fn_attr_get()` if the combined size of all the values is expected to be too large to be returned by `fn_attr_get()`.

There may be a relationship between the *ctx* argument supplied to `fn_attr_get_values()` and the `FN_valuelist_t` object it returns. For example, some implementations may store the context handle *ctx* within the `FN_valuelist_t` object for subsequent `fn_valuelist_next()` calls. In general, an `fn_ctx_handle_destroy(3XFN)` should not be invoked on *ctx* until the enumeration has terminated.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `FN_attribute_t(3XFN)`, `FN_attrvalue_t(3XFN)`, `FN_composite_name_t(3XFN)`, `FN_ctx_t(3XFN)`, `FN_identifier_t(3XFN)`, `FN_status_t(3XFN)`, `fn_attr_get(3XFN)`, `fn_ctx_handle_destroy(3XFN)`, `fn_ctx_list_names(3XFN)`, `xfn(3XFN)`, `xfn_attributes(3XFN)`, `xfn_status_codes(3XFN)`, `attributes(5)`

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.



<b>NAME</b>	FN_attribute_t, fn_attribute_create, fn_attribute_destroy, fn_attribute_copy, fn_attribute_assign, fn_attribute_identifier, fn_attribute_syntax, fn_attribute_valuecount, fn_attribute_first, fn_attribute_next, fn_attribute_add, fn_attribute_remove – an XFN attribute
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_attribute_t *fn_attribute_create(constFN_identifier_t     *attribute_id, const FN_identifier_t *attribute_syntax);  void fn_attribute_destroy(FN_attribute_t *attr);  FN_attribute_t *fn_attribute_copy(constFN_attribute_t *attr);  FN_attribute_t *fn_attribute_assign(FN_attribute_t *dst, const     FN_attribute_t *src);  const FN_identifier_t     *fn_attribute_identifier(constFN_attribute_t *attr);  const FN_identifier_t *fn_attribute_syntax(constFN_attribute_t     *attr);  unsigned int fn_attribute_valuecount(constFN_attribute_t *attr);  const FN_attrvalue_t *fn_attribute_first(constFN_attribute_t     *attr, void **iter_pos);  const FN_attrvalue_t *fn_attribute_next(constFN_attribute_t *attr,     void **iter_pos);  int fn_attribute_add(FN_attribute_t *attr, const FN_attrvalue_t     *attribute_value, unsigned int exclusive);  int fn_attribute_remove(FN_attribute_t *attr, const FN_attrvalue_t     *attribute_value);</pre>
<b>DESCRIPTION</b>	<p>An attribute has an attribute identifier, a syntax, and a set of distinct values. Each value is a sequence of octets. The operations associated with objects of type FN_attribute_t allow the construction, destruction, and manipulation of an attribute and its value set.</p> <p>The attribute identifier and its syntax are specified using an FN_identifier_t. fn_attribute_create() creates a new attribute object with the given identifier and syntax, and an empty set of values. fn_attribute_destroy() releases the storage associated with attr. fn_attribute_copy() returns a copy of the object pointed to by attr. fn_attribute_assign() makes a copy of the attribute object pointed to by src and assigns it to dst, releasing any old contents of dst. A pointer to the same object as dst is returned.</p> <p>fn_attribute_identifier() returns the attribute identifier of attr.  fn_attribute_syntax() returns the attribute syntax of attr.  fn_attribute_valuecount() returns the number of attribute values in attr.</p>

## FN\_attribute\_t(3XFN)

`fn_attribute_first()` and `fn_attribute_next()` are used to enumerate the values of an attribute. Enumeration of the values of an attribute may return the values in any order. `fn_attribute_first()` returns an attribute value from *attr* and sets the iteration marker *iter\_pos*. Subsequent calls to `fn_attribute_next()` returns the next attribute value identified by *iter\_pos* and advances *iter\_pos*. Adding or removing values from an attribute invalidates any iteration markers that the caller holds.

`fn_attribute_add()` adds a new value *attribute\_value* to *attr*. The operation succeeds (but no change is made) if *attribute\_value* is already in *attr* and *exclusive* is 0; the operation fails if *attribute\_value* is already in *attr* and *exclusive* is non-zero.

`fn_attribute_remove()` removes *attribute\_value* from *attr*. The operation succeeds even if *attribute\_value* is not amongst *attr*'s values.

### RETURN VALUES

`fn_attribute_first()` returns 0 if the attribute contains no values.  
`fn_attribute_next()` returns 0 if there are no more values to be returned in the attribute (as identified by the iteration marker) or if the iteration marker is invalid.

`fn_attribute_add()` and `fn_attribute_remove()` return 1 if the operation succeeds, 0 if it fails.

### USAGE

Manipulation of attributes using the operations described in this manual page does not affect their representation in the underlying naming system. Changes to attributes in the underlying naming system can only be effected through the use of the interfaces described in `xfn_attributes(3XFN)`.

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

### SEE ALSO

`FN_attrset_t(3XFN)`, `FN_attrvalue_t(3XFN)`, `FN_identifier_t(3XFN)`,  
`fn_attr_get(3XFN)`, `fn_attr_modify(3XFN)`, `xfn(3XFN)`,  
`xfn_attributes(3XFN)`, `attributes(5)`

### NOTES

The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

<b>NAME</b>	fn_attr_modify – modify specified attribute associated with name										
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxfn [ <i>library</i> ... ] #include &lt;xfn/xfn.h&gt;  int <b>fn_attr_modify</b>(FN_ctx_t *ctx, const FN_composite_name_t *name,     unsigned int <i>mod_op</i>, const FN_attribute_t *attr, unsigned int     <i>follow_link</i>, FN_status_t *status);</pre>										
<b>DESCRIPTION</b>	<p>This operation modifies according to <i>mod_op</i> the attribute <i>attr</i> associated with the object named <i>name</i> relative to <i>ctx</i>. If <i>name</i> is empty, the attribute associated with <i>ctx</i> is modified.</p> <p>The value of <i>follow_link</i> determines what happens when the terminal atomic part of <i>name</i> is bound to an XFN link. If <i>follow_link</i> is non-zero, such a link is followed, and the values of the attribute associated with the final named object are returned; if <i>follow_link</i> is zero, such a link is not followed. Any XFN links encountered before the terminal atomic name are always followed.</p> <p>The modification is made on the attribute identified by the attribute identifier of <i>attr</i>. The syntax and values of <i>attr</i> are used according to the modification operation.</p> <p>The modification operations are as follows:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">FN_ATTR_OP_ADD</td> <td>Add an attribute with given attribute identifier and set of values. If an attribute with this identifier already exists, replace the set of values with those in the given set. The set of values may be empty if the target naming system permits.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">FN_ATTR_OP_ADD_EXCLUSIVE</td> <td>Add an attribute with the given attribute identifier and set of values. The operation fails if an attribute with this identifier already exists. The set of values may be empty if the target naming system permits.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">FN_ATTR_OP_REMOVE</td> <td>Remove the attribute with the given attribute identifier and all of its values. The operation succeeds even if the attribute does not exist. The values of the attribute supplied with this operation are ignored.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">FN_ATTR_OP_ADD_VALUES</td> <td>Add the given values to those of the given attribute (resulting in the attribute having the union of its prior value set with the set given). Create the attribute if it does not exist already. The set of values may be empty if the target naming system permits.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">FN_ATTR_OP_REMOVE_VALUES</td> <td>Remove the given values from those of the given attribute (resulting in the attribute</td> </tr> </table>	FN_ATTR_OP_ADD	Add an attribute with given attribute identifier and set of values. If an attribute with this identifier already exists, replace the set of values with those in the given set. The set of values may be empty if the target naming system permits.	FN_ATTR_OP_ADD_EXCLUSIVE	Add an attribute with the given attribute identifier and set of values. The operation fails if an attribute with this identifier already exists. The set of values may be empty if the target naming system permits.	FN_ATTR_OP_REMOVE	Remove the attribute with the given attribute identifier and all of its values. The operation succeeds even if the attribute does not exist. The values of the attribute supplied with this operation are ignored.	FN_ATTR_OP_ADD_VALUES	Add the given values to those of the given attribute (resulting in the attribute having the union of its prior value set with the set given). Create the attribute if it does not exist already. The set of values may be empty if the target naming system permits.	FN_ATTR_OP_REMOVE_VALUES	Remove the given values from those of the given attribute (resulting in the attribute
FN_ATTR_OP_ADD	Add an attribute with given attribute identifier and set of values. If an attribute with this identifier already exists, replace the set of values with those in the given set. The set of values may be empty if the target naming system permits.										
FN_ATTR_OP_ADD_EXCLUSIVE	Add an attribute with the given attribute identifier and set of values. The operation fails if an attribute with this identifier already exists. The set of values may be empty if the target naming system permits.										
FN_ATTR_OP_REMOVE	Remove the attribute with the given attribute identifier and all of its values. The operation succeeds even if the attribute does not exist. The values of the attribute supplied with this operation are ignored.										
FN_ATTR_OP_ADD_VALUES	Add the given values to those of the given attribute (resulting in the attribute having the union of its prior value set with the set given). Create the attribute if it does not exist already. The set of values may be empty if the target naming system permits.										
FN_ATTR_OP_REMOVE_VALUES	Remove the given values from those of the given attribute (resulting in the attribute										

## fn\_attr\_modify(3XFN)

having the set difference of its prior value set and the set given). This succeeds even if some of the given values are not in the set of values that the attribute has. In naming systems that require an attribute to have at least one value, removing the last value will remove the attribute as well.

**RETURN VALUES** 1 Successful operation.  
0 Operation failed.

**ERRORS** `fn_attr_modify()` sets *status* as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `FN_attribute_t(3XFN)`, `FN_composite_name_t(3XFN)`, `FN_ctx_t(3XFN)`, `FN_status_t(3XFN)`, `fn_attr_multi_modify(3XFN)`, `xfn(3XFN)`, `xfn_attributes(3XFN)`, `xfn_status_codes(3XFN)`, `attributes(5)`

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

<b>NAME</b>	FN_attrmodlist_t, fn_attrmodlist_create, fn_attrmodlist_destroy, fn_attrmodlist_copy, fn_attrmodlist_assign, fn_attrmodlist_count, fn_attrmodlist_first, fn_attrmodlist_next, fn_attrmodlist_add – a list of attribute modifications
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_attrmodlist_t *fn_attrmodlist_create(void);  void fn_attrmodlist_destroy(FN_attrmodlist_t *modlist);  FN_attrmodlist_t *fn_attrmodlist_copy(const FN_attrmodlist_t     *modlist);  FN_attrmodlist_t *fn_attrmodlist_assign(FN_attrmodlist_t *dst,     const FN_attrmodlist_t *src);  unsigned int fn_attrmodlist_count(const FN_attrmodlist_t *modlist);  const FN_attribute_t *fn_attrmodlist_first(const     FN_attrmodlist_t *modlist, void **iter_pos, unsigned int     *first_mod_op);  const FN_attribute_t *fn_attrmodlist_next(const FN_attrmodlist_t     *modlist, void **iter_pos, unsigned int *mod_op);  int fn_attrmodlist_add(FN_attrmodlist_t *modlist, unsigned int     mod_op, const FN_attribute_t *attr);</pre>
<b>DESCRIPTION</b>	<p>An attribute modification list allows for multiple modification operations to be made on the attributes associated with a single named object. It is used in the <code>fn_attr_multi_modify(3XFN)</code> operation.</p> <p>An attribute modification list is a list of attribute modification specifiers. An attribute modification specifier consists of an attribute object and an operation specifier. The attribute's identifier indicates the attribute that is to be operated upon. The attribute's values are used in a manner depending on the operation. The operation specifier is an unsigned int that must have one of the values:</p> <pre>FN_ATTR_OP_ADD FN_ATTR_OP_ADD_EXCLUSIVE FN_ATTR_OP_REMOVE FN_ATTR_OP_ADD_VALUES</pre> <p>or</p> <pre>FN_ATTR_OP_REMOVE_VALUES</pre> <p>(See <code>fn_attr_modify(3XFN)</code> for detailed descriptions of these specifiers.) The operations are to be performed in the order in which they appear in the modification list.</p>

## FN\_attrmodlist\_t(3XFN)

`fn_attrmodlist_create()` creates an empty attribute modification list.  
`fn_attrmodlist_destroy()` releases the storage associated with *modlist*.  
`fn_attrmodlist_copy()` returns a copy of the attribute modification list *modlist*.  
`fn_attrmodlist_assign()` makes a copy of *src* and assigns it to *dst*, releasing any old contents of *dst*. It returns a pointer to the same object as *dst*.

`fn_attrmodlist_count()` returns the number attribute modification items in the attribute modification list.

The iterators `fn_attrmodlist_first()` and `fn_attrmodlist_next()` return a handle to the attribute part of the modification and return the operation specifier part through an `unsigned int *` parameter. `fn_attrmodlist_first()` returns the attribute of the first modification item from *modlist* and sets *mod\_op* to be the code of the modification operation of that item; *iter\_pos* is set after the first modification item.

`fn_attrmodlist_next()` returns the attribute of the next modification item from *modlist* after *iter\_pos* and advances *iter\_pos*; *mod\_op* is set to the code of the modification operation of that item. The order of the items returned during an enumeration is the same as the order by which the items were added to the modification list.

`fn_attrmodlist_add()` adds a new item consisting of the given modification operation code *mod\_op* and attribute *attr* to the end of the modification list *modlist*. *attr*'s identifier indicates the attribute that is to be operated upon. *attr*'s values are used in a manner depending on the operation.

### RETURN VALUES

`fn_attrmodlist_first()` returns 0 if the modification list is empty.  
`fn_attrmodlist_next()` returns 0 if there are no more items on the modification list to be enumerated or if the iteration marker is invalid.

`fn_attrmodlist_add()` returns 1 if the operation succeeds, 0 if the operation fails.

### USAGE

Manipulation of attributes using the operations described in this manual page does not affect their representation in the underlying naming system. Changes to attributes in the underlying naming system can only be effected through the use of the interfaces described in `xfn_attributes(3XFN)`.

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

### SEE ALSO

`FN_attribute_t(3XFN)`, `FN_attrset_t(3XFN)`, `FN_identifier_t(3XFN)`,  
`fn_attr_modify(3XFN)`, `fn_attr_multi_modify(3XFN)`, `xfn(3XFN)`,  
`xfn_attributes(3XFN)`, `attributes(5)`

**NOTES** | The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_attr\_multi\_get(3XFN)

<b>NAME</b>	fn_attr_multi_get, FN_multigetlist_t, fn_multigetlist_next, fn_multigetlist_destroy – return multiple attributes associated with named object
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_multigetlist_t *fn_attr_multi_get(FN_ctx_t *ctx, const     FN_composite_name_t *name, const FN_attrset_t *attr_ids,     unsigned int follow_link, FN_status_t *status);  FN_attribute_t *fn_multigetlist_next(FN_multigetlist_t *ml,     FN_status_t *status);  void fn_multigetlist_destroy(FN_multigetlist_t *ml, FN_status_t     *status);</pre>
<b>DESCRIPTION</b>	<p>This set of operations returns one or more attributes associated with the object named by <i>name</i> relative to the context <i>ctx</i>. If <i>name</i> is empty, the attributes associated with <i>ctx</i> are returned.</p> <p>The value of <i>follow_link</i> determines what happens when the terminal atomic part of <i>name</i> is bound to an XFN link. If <i>follow_link</i> is non-zero, such a link is followed, and the values of the attribute associated with the final named object are returned; if <i>follow_link</i> is zero, such a link is not followed. Any XFN links encountered before the terminal atomic name are always followed.</p> <p>The attributes returned are those specified in <i>attr_ids</i>. If the value of <i>attr_ids</i> is 0, all attributes associated with the named object are returned. Any attribute values in <i>attr_ids</i> provided by the caller are ignored; only the attribute identifiers are relevant for this operation. Each attribute (identifier, syntax, values) is returned one at a time using an enumeration scheme similar to that for listing a context.</p> <p>fn_attr_multi_get() initiates the enumeration process. It returns a handle to an FN_multigetlist_t object that can be used for the enumeration.</p> <p>The operation fn_multigetlist_next() returns a new FN_attribute_t object containing the next attribute (identifiers, syntaxes, and values) requested and updates <i>ml</i> to indicate the state of the enumeration.</p> <p>The operation fn_multigetlist_destroy() releases the resources used during the enumeration. It may be invoked before the enumeration has completed to terminate the enumeration.</p>
<b>RETURN VALUES</b>	<p>fn_attr_multi_get() returns a pointer to an FN_multigetlist_t object if the enumeration has been initiated successfully; a NULL pointer (0) is returned if it failed.</p> <p>fn_multigetlist_next() returns a pointer to an FN_attribute_t object if an attribute was returned, a NULL pointer (0) if no attribute was returned.</p> <p>In the case of a failure, these operations set <i>status</i> to indicate the nature of the failure.</p>
<b>ERRORS</b>	Each call to fn_multigetlist_next() sets <i>status</i> as follows:



## fn\_attr\_multi\_get(3XFN)

FN_SUCCESS	If an attribute was returned, there are more attributes to be enumerated. If no attribute was returned, the enumeration has completed successfully.
FN_E_ATTR_NO_PERMISSION	The caller did not have permission to read this attribute.
FN_E_INSUFFICIENT_RESOURCES	Insufficient resources are available to return the attribute's values.
FN_E_INVALID_ATTR_IDENTIFIER	This attribute identifier was not in a format acceptable to the naming system, or its contents was not valid for the format specified for the identifier.
FN_E_INVALID_ENUM_HANDLE	(No attribute should be returned with this status code). The given enumeration handle is not valid. Possible reasons could be that the handle was from another enumeration, or the object being processed no longer accepts the handle (due to such events as handle expiration or updates to the object's attribute set).
FN_E_NO_SUCH_ATTRIBUTE	The object did not have an attribute with the given identifier.
FN_E_PARTIAL_RESULT	(No attribute should be returned with this status code). The enumeration is not yet complete but cannot be continued.

For `FN_E_ATTR_NO_PERMISSION`, `FN_E_INVALID_ATTR_IDENTIFIER`, `FN_E_INSUFFICIENT_RESOURCES`, or `FN_E_NO_SUCH_ATTRIBUTE`, the returned attribute contains only the attribute identifier (no value or syntax). For these four status codes and `FN_SUCCESS` (when an attribute was returned), `fn_multigetlist_next()` can be called again to return another attribute. All other status codes indicate that no more attributes can be returned by `fn_multigetlist_next()`.

Other status codes, such as `FN_E_COMMUNICATION_FAILURE`, are also possible, in which case, no attribute is returned. In such cases, *status* is set as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`.

**USAGE** Implementations are not required to return all attributes requested by *attr\_ids*. Some may choose to return only the attributes found successfully, followed by a status of `FN_E_PARTIAL_RESULT`; such implementations may not necessarily return attributes identifying those that could not be read. Implementations are not required to return the attributes in any order.

## fn\_attr\_multi\_get(3XFN)

There may be a relationship between the *ctx* argument supplied to `fn_attr_multi_get()` and the `FN_multigetlist_t` object it returns. For example, some implementations may store the context handle *ctx* within the `FN_multigetlist_t` object for subsequent `fn_multigetlist_next()` calls. In general, a `fn_ctx_handle_destroy()` should not be invoked on *ctx* until the enumeration has terminated.

**EXAMPLES**     **EXAMPLE 1** A sample program displaying how to use `fn_attr_multi_get()` function.

The following code fragment illustrates to obtain all attributes associated with a given name using the `fn_attr_multi_get()` operations.

```
/* list all attributes associated with given name */
extern FN_string_t *input_string;
FN_ctx_t *ctx;
FN_composite_name_t *target_name = fn_composite_name_from_string(input_string);
FN_multigetlist_t *ml;
FN_status_t *status = fn_status_create();
FN_attribute_t *attr;
int done = 0;
ctx = fn_ctx_handle_from_initial(status);
/* error checking on 'status' */
/* attr_ids == 0 indicates all attributes are to be returned */
if ((ml=fn_attr_multi_get(ctx, target_name, 0, status)) == 0) {
    /* report 'status' and exit */
}
while ((attr=fn_multigetlist_next(ml, status)) && !done) {
    switch (fn_status_code(status)) {
        case FN_SUCCESS:
            /* do something with 'attr' */
            break;
        case FN_E_ATTR_NO_PERMISSION:
        case FN_E_ATTR_INVALID_ATTR_IDENTIFIER:
        case FN_E_NO_SUCH_ATTRIBUTE:
            /* report error using identifier in 'attr' */
            break;
        default:
            /* other error handling */
            done = 1;
    }
    if (attr)
        fn_attribute_destroy(attr);
}
/* check 'status' for reason for end of enumeration and report if necessary */
/* clean up */
fn_multigetlist_destroy(ml, status);
/* report 'status' */
```

**ATTRIBUTES**     See `attributes(5)` for descriptions of the following attributes:

fn\_attr\_multi\_get(3XFN)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_attribute\_t(3XFN), FN\_attrset\_t(3XFN), FN\_composite\_name\_t(3XFN), FN\_ctx\_t(3XFN), FN\_identifier\_t(3XFN), FN\_status\_t(3XFN), fn\_attr\_get(3XFN), fn\_ctx\_handle\_destroy(3XFN), fn\_ctx\_list\_names(3XFN), xfn(3XFN), xfn\_attributes(3XFN), xfn\_status\_codes(3XFN), attributes(5)

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_attr\_multi\_modify(3XFN)

<b>NAME</b>	fn_attr_multi_modify – modify multiple attributes associated with named object				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  int <b>fn_attr_multi_modify</b>(FN_ctx_t *ctx, const FN_composite_name_t     *name, const FN_attrmodlist_t *mods, unsigned int follow_link,     FN_attrmodlist_t **unexecuted_mods, FN_status_t *status);</pre>				
<b>DESCRIPTION</b>	<p>This operation modifies the attributes associated with the object named <i>name</i> relative to <i>ctx</i>. If <i>name</i> is empty, the attributes associated with <i>ctx</i> are modified.</p> <p>The value of <i>follow_link</i> determines what happens when the terminal atomic part of <i>name</i> is bound to an XFN link. If <i>follow_link</i> is non-zero, such a link is followed, and the values of the attribute associated with the final named object are returned; if <i>follow_link</i> is zero, such a link is not followed. Any XFN links encountered before the terminal</p> <p>In the <i>mods</i> parameter, the caller specifies a sequence of modifications that are to be done in order on the attributes. Each modification in the sequence specifies a modification operation code (see <code>fn_attr_modify(3XFN)</code>) and an attribute on which to operate.</p> <p>The <code>FN_attrmodlist_t</code> type is described in <code>FN_attrmodlist_t(3XFN)</code>.</p>				
<b>RETURN VALUES</b>	<code>fn_attr_multi_modify()</code> returns 1 if all the modification operations were performed successfully. The function returns 0 if any error occurs. If the operation fails, <i>status</i> and <i>unexecuted_mods</i> are set as described below.				
<b>ERRORS</b>	If an error is encountered while performing the list of modifications, <i>status</i> indicates the type of error and <i>unexecuted_mods</i> is set to a list of unexecuted modifications. The contents of <i>unexecuted_mods</i> do not share any state with <i>mods</i> ; items in <i>unexecuted_mods</i> are copies of items in <i>mods</i> and appear in the same order in which they were originally supplied in <i>mods</i> . The first operation in <i>unexecuted_mods</i> is the first one that failed and the code in <i>status</i> applies to this modification operation in particular. If <i>status</i> indicates failure and a NULL pointer (0) is returned in <i>unexecuted_mods</i> , that indicates no modifications were executed.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>FN_attrmodlist_t(3XFN)</code> , <code>FN_composite_name_t(3XFN)</code> , <code>FN_ctx_t(3XFN)</code> , <code>FN_status_t(3XFN)</code> , <code>fn_attr_modify(3XFN)</code> , <code>xfn(3XFN)</code> , <code>xfn_attributes(3XFN)</code> , <code>xfn_status_codes(3XFN)</code> , <code>attributes(5)</code>				

fn\_attr\_multi\_modify(3XFN)

**NOTES** | The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_attr\_search(3XFN)

<b>NAME</b>	fn_attr_search, FN_searchlist_t, fn_searchlist_next, fn_searchlist_destroy – search for the atomic name of objects with the specified attributes in a single context
<b>SYNOPSIS</b>	<pre>#include &lt;xfn/xfn.h&gt;  FN_searchlist_t *fn_attr_search(FN_ctx_t *ctx, const     FN_composite_name_t *name, const FN_attrset_t *match_attrs,     unsigned int return_ref, const FN_attrset_t *return_attr_ids,     FN_status_t *status);  FN_string_t *fn_searchlist_next(FN_searchlist_t *sl, FN_ref_t     **returned_ref, FN_attrset_t **returned_attrs, FN_status_t *status);  void fn_searchlist_destroy(FN_searchlist_t *sl);</pre>
<b>DESCRIPTION</b>	<p>This set of operations is used to enumerate names of objects bound in the target context named <i>name</i> relative to the context <i>ctx</i> with attributes whose values match all those specified by <i>match_attrs</i>.</p> <p>The attributes specified by <i>match_attrs</i> form a conjunctive AND expression against which the attributes of each named object in the target context are evaluated. For multi-valued attributes, the list order of values is ignored and attribute values not specified in <i>match_attrs</i> are ignored. If no value is specified for an attribute in <i>match_attrs</i>, the presence of the attribute is tested. If the value of <i>match_attrs</i> is 0, all names in the target context are enumerated.</p> <p>If a non-zero value of <i>return_ref</i> is passed to <code>fn_attr_search()</code>, the reference bound to the name is returned in the <i>returned_ref</i> argument to <code>fn_searchlist_next()</code>.</p> <p>Attribute identifiers and values associated with named objects that satisfy <i>match_attrs</i> may be returned by <code>fn_searchlist_next()</code>. The attributes returned are those listed in the <i>return_attr_ids</i> argument to <code>fn_attr_search()</code>. If the value of <i>return_attr_ids</i> is 0, all attributes are returned. If <i>return_attr_ids</i> is an empty <code>FN_attrset_t(3XFN)</code> object, no attributes are returned. Any attribute values in <i>return_attr_ids</i> are ignored; only the attribute identifiers are relevant for <i>return_attr_ids</i>.</p> <p>The call to <code>fn_attr_search()</code> initiates the enumeration process. It returns a handle to an <code>FN_searchlist_t</code> object that is used to enumerate the names of the objects whose attributes match the attributes specified by <i>match_attrs</i>.</p> <p>The operation <code>fn_searchlist_next()</code> returns the next name in the enumeration identified by the <i>sl</i>. The reference of the name is returned in <i>returned_ref</i> if <i>return_ref</i> was set in the call to <code>fn_attr_search()</code>. The attributes specified by <i>return_attr_ids</i> are returned in <i>returned_attrs</i>. <code>fn_searchlist_next()</code> also updates <i>sl</i> to indicate the state of the enumeration. Successive calls to <code>fn_searchlist_next()</code> using <i>sl</i> return successive names, and optionally, references and attributes, in the enumeration; these calls further update the state of the enumeration.</p> <p><code>fn_searchlist_destroy()</code> releases resources used during the enumeration. This can be invoked at any time to terminate the enumeration.</p>

**RETURN VALUES**

fn\_attr\_search() does not follow XFN links that are bound in the target context.

fn\_attr\_search() returns a pointer to an FN\_searchlist\_t object if the enumeration is successfully initiated; it returns a NULL pointer if the enumeration cannot be initiated or if no named object with attributes whose values match those specified in match\_attrs is found.

fn\_searchlist\_next() returns a pointer to an FN\_string\_t(3XFN) object; it returns a NULL pointer if no more names can be returned in the enumeration. If returned\_ref is a NULL pointer, or if the return\_ref parameter to fn\_attr\_search was 0, no reference is returned; otherwise, returned\_ref contains the reference bound to the name. If returned\_attrs is a NULL pointer, no attributes are returned; otherwise, returned\_attrs contains the attributes associated with the named object, as specified by the return\_attr\_ids parameter to fn\_attr\_search().

In the case of a failure, these operations return in the status argument a code indicating the nature of the failure.

**ERRORS**

fn\_attr\_search() returns a NULL pointer if the enumeration could not be initiated. The status argument is set in the following way:

FN_SUCCESS	A named object could not be found whose attributes satisfied the implied filter of equality and conjunction.
FN_E_ATTR_NO_PERMISSION	The caller did not have permission to read one or more of the specified attributes.
FN_E_INVALID_ATTR_VALUE	A value type in the specified attributes did not match the syntax of the attribute against which it was being evaluated.

Other status codes are possible as described in FN\_status\_t(3XFN) and xfn\_status\_codes(3XFN).

Each successful call to fn\_searchlist\_next() returns a name and, optionally, the reference and requested attributes. status is set in the following way:

FN_SUCCESS	All requested attributes were returned successfully with the name.
FN_E_ATTR_NO_PERMISSION	The caller did not have permission to read one or more of the requested attributes.
FN_E_INVALID_ATTR_IDENTIFIER	A requested attribute identifier was not in a format acceptable to the naming system, or its contents was not valid for the format specified.
FN_E_NO_SUCH_ATTRIBUTE	The named object did not have one of the requested attributes.

## fn\_attr\_search(3XFN)

FN_E_INSUFFICIENT_RESOURCES	Insufficient resources are available to return all the requested attributes and their values.
FN_E_ATTR_NO_PERMISSION	
FN_E_INVALID_ATTR_IDENTIFIER	
FN_E_NO_SUCH_ATTRIBUTE	
FN_E_INSUFFICIENT_RESOURCES	These indicate that some of the requested attributes may have been returned in <i>returned_attrs</i> but one or more of them could not be returned. Use <code>fn_attr_get(3XFN)</code> or <code>fn_attr_multi_get(3XFN)</code> to discover why these attributes could not be returned.

`fn_searchlist_next()` returns a NULL pointer if no more names can be returned. The status argument is set in the following way:

FN_SUCCESS	The search has completed successfully.
FN_E_PARTIAL_RESULT	The enumeration is not yet complete but cannot be continued.
FN_E_ATTR_NO_PERMISSION	The caller did not have permission to read one or more of the specified attributes.
FN_E_INVALID_ENUM_HANDLE	The supplied enumeration handle was not valid. Possible reasons could be that the handle was from another enumeration, or the context being enumerated no longer accepts the handle (due to such events as handle expiration or updates to the context).

Other status codes are possible as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`.

**USAGE** The names enumerated using `fn_searchlist_next()` are not ordered in any way. Furthermore, there is no guarantee that any two series of enumerations on the same context with identical *match\_attrs* will return the names in the same order.

**EXAMPLES** **EXAMPLE 1** A sample program of displaying how to use `fn_attr_search()` function.

The following code fragment illustrates how the `fn_attr_search()` operation may be used. The code consists of three parts: preparing the arguments for the search, performing the search, and cleaning up.

The first part involves getting the name of the context to start the search and constructing the set of attributes that named objects in the context must satisfy. This is done in the declarations part of the code and by the routine `get_search_query`.



**EXAMPLE 1** A sample program of displaying how to use `fn_attr_search()` function.  
(Continued)

The next part involves doing the search and enumerating the results of the search. This is done by first getting a context handle to the Initial Context, and then passing that handle along with the name of the target context and matching attributes to `fn_attr_search()`. This particular call to `fn_attr_search()` is requesting that no reference be returned (by passing in 0 for `return_ref`), and that all attributes associated with the named object be returned (by passing in 0 as the `return_attr_ids` argument). If successful, `fn_attr_search()` returns `sl`, a handle for enumerating the results of the search. The results of the search are enumerated using calls to `fn_searchlist_next()`, which returns the name of the object and the attributes associated with the named object in `returned_attrs`.

The last part of the code involves cleaning up the resources used during the search and enumeration. The call to `fn_searchlist_destroy()` releases resources reserved for this enumeration. The other calls release the context handle, name, attribute set, and status objects created earlier.

```

/* Declarations */
FN_ctx_t *ctx;
FN_searchlist_t *sl;
FN_string_t *name;
FN_attrset_t *returned_attrs;
FN_status_t *status = fn_status_create();
FN_composite_name_t *target_name = get_name_from_user_input();
FN_attrset_t *match_attrs = get_search_query();
/* Get context handle to Initial Context */
ctx = fn_ctx_handle_from_initial(status);
/* error checking on 'status' */
/* Initiate search */
if ((sl=fn_attr_search(ctx, target_name, match_attrs,
/* no reference */ 0, /* return all attrs */ 0, status)) == 0) {
/* report 'status', cleanup, and exit */
}
/* Enumerate names and attributes requested */
while (name=fn_searchlist_next(sl, 0, &returned_attrs, status)) {
/* do something with 'name' and 'returned_attrs'*/
fn_string_destroy(name);
fn_attrset_destroy(returned_attrs);
}
/* check 'status' for reason for end of enumeration */
/* Clean up */
fn_searchlist_destroy(sl); /* Free resources of 'sl' */
fn_status_destroy(status);
fn_attrset_destroy(match_attrs);
fn_ctx_handle_destroy(ctx);
fn_composite_name_destroy(target_name);
/*
* Procedure for constructing attribute set containing
* attributes to be matched:
* "zip_code" attribute value is "02158"
* AND "employed" attribute is present.
*/

```

## fn\_attr\_search(3XFN)

**EXAMPLE 1** A sample program of displaying how to use `fn_attr_search()` function.  
(Continued)

```
FN_attrset_t *
get_search_query()
{
    /* Zip code and employed attribute identifier, syntax */
    extern FN_attribute_t      *attr_zip_code;
    extern FN_attribute_t      *attr_employed;
    FN_attribute_t *zip_code = fn_attribute_copy(attr_zip_code);
    FN_attr_value_t zc_value = {5, "02158"};
    FN_attrset_t *match_attrs = fn_attrset_create();
    fn_attribute_add(zip_code, &zc_value, 0);
    fn_attrset_add(match_attrs, zip_code, 0);
    fn_attrset_add(match_attrs, attr_employed, 0);
    return (match_attrs);
}
```

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `FN_attribute_t(3XFN)`, `FN_attrset_t(3XFN)`, `FN_attrvalue_t(3XFN)`,  
`FN_composite_name_t(3XFN)`, `FN_ctx_t(3XFN)`, `FN_status_t(3XFN)`,  
`FN_string_t(3XFN)`, `fn_attr_ext_search(3XFN)`, `fn_attr_get(3XFN)`,  
`fn_attr_multi_get(3XFN)`, `fn_ctx_list_names(3XFN)`,  
`xfn_status_codes(3XFN)`, `attributes(5)`

<b>NAME</b>	FN_attrset_t, fn_attrset_create, fn_attrset_destroy, fn_attrset_copy, fn_attrset_assign, fn_attrset_get, fn_attrset_count, fn_attrset_first, fn_attrset_next, fn_attrset_add, fn_attrset_remove – a set of XFN attributes
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_attrset_t *fn_attrset_create(void);  void fn_attrset_destroy(FN_attrset_t *aset);  FN_attrset_t *fn_attrset_copy(constFN_attrset_t *aset);  FN_attrset_t *fn_attrset_assign(FN_attrset_t *dst, const     FN_attrset_t *src);  const FN_attribute_t *fn_attrset_get(constconst FN_attrset_t     *aset, const FN_identifier_t *attr_id);  unsigned int fn_attrset_count(constFN_attrset_t *aset);  const FN_attribute_t *fn_attrset_first(constFN_attrset_t *aset,     void **iter_pos);  const FN_attribute_t *fn_attrset_next(constFN_attrset_t *aset,     void **iter_pos);  int fn_attrset_add(FN_attrset_t *aset, const FN_attribute_t *attr,     unsigned int exclusive);  int fn_attrset_remove(FN_attrset_t *aset, const FN_identifier_t     *attr_id);</pre>
<b>DESCRIPTION</b>	<p>An attribute set is a set of attribute objects with distinct identifiers. The <code>fn_attr_multi_get(3XFN)</code> operation takes an attribute set as parameter and returns an attribute set. The <code>fn_attr_get_ids(3XFN)</code> operation returns an attribute set containing the identifiers of the attributes.</p> <p>Attribute sets are represented by the type <code>FN_attrset_t</code>. The following operations are defined for manipulating attribute sets.</p> <p><code>fn_attrset_create()</code> creates an empty attribute set. <code>fn_attrset_destroy()</code> releases the storage associated with the attribute set <code>aset</code>. <code>fn_attrset_copy()</code> returns a copy of the attribute set <code>aset</code>. <code>fn_attrset_assign()</code> makes a copy of the attribute set <code>src</code> and assigns it to <code>dst</code>, releasing any old contents of <code>dst</code>. A pointer to the same object as <code>dst</code> is returned.</p> <p><code>fn_attrset_get()</code> returns the attribute with the given identifier <code>attr_id</code> from <code>aset</code>. <code>fn_attrset_count()</code> returns the number attributes found in the attribute set <code>aset</code>.</p> <p><code>fn_attrset_first()</code> and <code>fn_attrset_next()</code> are functions that can be used to return an enumeration of all the attributes in an attribute set. The attributes are not ordered in any way. There is no guaranteed relation between the order in which items are added to an attribute set and the order of the enumeration. The specification does</p>

## FN\_attrset\_t(3XFN)

guarantee that any two enumerations will return the members in the same order, provided that no `fn_attrset_add()` or `fn_attrset_remove()` operation was performed on the object in between or during the two enumerations.

`fn_attrset_first()` returns the first attribute from the set and sets `iter_pos` after the first attribute. `fn_attrset_next()` returns the attribute following `iter_pos` and advances `iter_pos`.

`fn_attrset_add()` adds the attribute `attr` to the attribute set `aset`, replacing the attribute's values if the identifier of `attr` is not distinct in `aset` and `exclusive` is 0. If `exclusive` is non-zero and the identifier of `attr` is not distinct in `aset`, the operation fails.

`fn_attrset_remove()` removes the attribute with the identifier `attr_id` from `aset`. The operation succeeds even if no such attribute occurs in `aset`.

### RETURN VALUES

`fn_attrset_first()` returns 0 if the attribute set is empty. `fn_attrset_next()` returns 0 if there are no more attributes in the set.

`fn_attrset_add()` and `fn_attrset_remove()` return 1 if the operation succeeds, and 0 if the operation fails.

### USAGE

Manipulation of attributes using the operations described in this manual page does not affect their representation in the underlying naming system. Changes to attributes in the underlying naming system can only be effected through the use of the interfaces described in `xfn_attributes(3XFN)`.

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

### SEE ALSO

`FN_attribute_t(3XFN)`, `FN_attrvalue_t(3XFN)`, `FN_identifier_t(3XFN)`, `fn_attr_get_ids(3XFN)`, `fn_attr_multi_get(3XFN)`, `xfn(3XFN)`, `xfn_attributes(3XFN)`, `attributes(5)`

### NOTES

The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

<b>NAME</b>	FN_attrvalue_t – an XFN attribute value
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxfn [ <i>library</i> ... ] #include &lt;xfn/xfn.h&gt;</pre>
<b>DESCRIPTION</b>	<p>The type FN_attrvalue_t is used to represent the contents of a single attribute value, within an attribute of type FN_attribute_t.</p> <p>The representation of this structure is defined by XFN as follows:</p> <pre>typedef struct { size_t length; void *contents; } FN_attrvalue_t;</pre>
<b>SEE ALSO</b>	FN_attribute_t(3XFN), fn_attr_get_values(3XFN), xfn(3XFN)

## FN\_composite\_name\_t(3XFN)

<b>NAME</b>	FN_composite_name_t, fn_composite_name_create, fn_composite_name_destroy, fn_composite_name_from_str, fn_composite_name_from_string, fn_string_from_composite_name, fn_composite_name_copy, fn_composite_name_assign, fn_composite_name_is_empty, fn_composite_name_count, fn_composite_name_first, fn_composite_name_next, fn_composite_name_prev, fn_composite_name_last, fn_composite_name_prefix, fn_composite_name_suffix, fn_composite_name_is_equal, fn_composite_name_is_prefix, fn_composite_name_is_suffix, fn_composite_name_prepend_comp, fn_composite_name_append_comp, fn_composite_name_insert_comp, fn_composite_name_delete_comp, fn_composite_name_prepend_name, fn_composite_name_append_name, fn_composite_name_insert_name – a sequence of component names spanning multiple naming systems
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_composite_name_t *fn_composite_name_create(void); void fn_composite_name_destroy(FN_composite_name_t *name); FN_composite_name_t *fn_composite_name_from_str(const unsigned char *cstr); FN_composite_name_t *fn_composite_name_from_string(const FN_string_t *str); FN_string_t *fn_string_from_composite_name(const FN_composite_name_t *name, unsigned int *status); FN_composite_name_t *fn_composite_name_copy(const FN_composite_name_t *name); FN_composite_name_t *fn_composite_name_assign(FN_composite_name_t *dst, const FN_composite_name_t *src); int fn_composite_name_is_empty(const FN_composite_name_t *name); unsigned int fn_composite_name_count(const FN_composite_name_t *name); const FN_string_t *fn_composite_name_first(const FN_composite_name_t *name, void **iter_pos); const FN_string_t *fn_composite_name_next(const FN_composite_name_t *name, void **iter_pos); const FN_string_t *fn_composite_name_prev(const FN_composite_name_t *name, void **iter_pos); const FN_string_t *fn_composite_name_last(const FN_composite_name_t *name, void **iter_pos);</pre>

## FN\_composite\_name\_t(3XFN)

```

FN_composite_name_t *fn_composite_name_prefix(const
    FN_composite_name_t *name, const void *iter_pos);

FN_composite_name_t *fn_composite_name_suffix(const
    FN_composite_name_t *name, const void *iter_pos);

int fn_composite_name_is_equal(const FN_composite_name_t *name,
    const FN_composite_name_t *name2, unsigned int *status);

int fn_composite_name_is_prefix(const FN_composite_name_t *name,
    const FN_composite_name_t *prefix, void **iter_pos, unsigned int
    *status);

int fn_composite_name_is_suffix(const FN_composite_name_t *name,
    const FN_composite_name_t *suffix, void **iter_pos, unsigned int
    *status);

int fn_composite_name_prepend_comp(FN_composite_name_t *name,
    const FN_string_t *newcomp);

int fn_composite_name_append_comp(FN_composite_name_t *name,
    const FN_string_t *newcomp);

int fn_composite_name_insert_comp(FN_composite_name_t *name, void
    **iter_pos, const FN_string_t *newcomp);

int fn_composite_name_delete_comp(FN_composite_name_t *name, void
    **iter_pos);

int fn_composite_name_prepend_name(FN_composite_name_t *name,
    const FN_composite_name_t *newcomps);

int fn_composite_name_append_name(FN_composite_name_t *name,
    const FN_composite_name_t *newcomps);

int fn_composite_name_insert_name(FN_composite_name_t *name, void
    **iter_pos, const FN_composite_name_t *newcomps);

```

### DESCRIPTION

A composite name is represented by an object of type `FN_composite_name_t`. Each component is a string name, of type `FN_string_t`, from the namespace of a single naming system. It may be an atomic name or a compound name in that namespace.

`fn_composite_name_create` creates an `FN_composite_name_t` object with zero components. Components may be subsequently added to the composite name using the modify operations described below. `fn_composite_name_destroy` releases any storage associated with the given `FN_composite_name_t` handle.

`fn_composite_name_from_str()` creates an `FN_composite_name_t` from the given null-terminated string based on the code set of the current locale setting, using the XFN composite name syntax. `fn_composite_name_from_string()` creates an `FN_composite_name_t` from the string `str` using the XFN composite name syntax. `fn_string_from_composite_name()` returns the standard string form of the given composite name, by concatenating the components of the composite name in a left to right order, each separated by the XFN component separator.

## FN\_composite\_name\_t(3XFN)

`fn_composite_name_copy()` returns a copy of the given composite name object. `fn_composite_name_assign()` makes a copy of the composite name object pointed to by *src* and assigns it to *dst*, releasing any old contents of *dst*. A pointer to the same object as *dst* is returned.

`fn_composite_name_is_empty()` returns 1 if the given composite name is an empty composite name (that is, it consists of a single, empty component name); otherwise, it returns 0. `fn_composite_name_count()` returns the number of components in the given composite name.

The iteration scheme is based on the exchange of an opaque `void *` argument, *iter\_pos*, that serves to record the position of the iteration in the sequence. Conceptually, *iter\_pos* records a position between two successive components (or at one of the extreme ends of the sequence).

The function `fn_composite_name_first()` returns a handle to the `FN_string_t` that is the first component in the name, and sets *iter\_pos* to indicate the position immediately following the first component. It returns 0 if the name has no components. Thereafter, successive calls of the `fn_composite_name_next()` function return pointers to the component following the iteration marker, and advance the iteration marker. If the iteration marker is at the end of the sequence, `fn_composite_name_next()` returns 0. Similarly, `fn_composite_name_prev()` returns the component preceding the iteration pointer and moves the marker back one component. If the marker is already at the beginning of the sequence, `fn_composite_name_prev()` returns 0. The function `fn_composite_name_last()` returns a pointer to the last component of the name and sets the iteration marker immediately preceding this component (so that subsequent calls to `fn_composite_name_prev()` can be used to step through leading components of the name).

The `fn_composite_name_suffix()` function returns a composite name consisting of a copy of those components following the supplied iteration marker. The method `fn_composite_name_prefix()` returns a composite name consisting of those components that precede the iteration marker. Using these functions with an iteration marker that was not initialized using `fn_composite_name_first()`, `fn_composite_name_last()`, `fn_composite_name_is_prefix()`, or `fn_composite_name_is_suffix()` yields undefined and generally undesirable behavior.

The functions `fn_composite_name_is_equal()`, `fn_composite_name_is_prefix()`, and `fn_composite_name_is_suffix()` test for equality between composite names or between parts of composite names. For these functions, equality is defined as exact string equality, not name equivalence. A name's syntactic property, such as case-insensitivity, is not taken into account by these functions.

The function `fn_composite_name_is_prefix()` tests if one composite name is a prefix of another. If so, it returns 1 and sets the iteration marker immediately following the prefix. (For example, a subsequent call to



`fn_composite_name_suffix()` will return the remainder of the name.) Otherwise, it returns 0 and the value of the iteration marker is undefined. The function `fn_composite_name_is_suffix()` is similar. It tests if one composite name is a suffix of another. If so, it returns 1 and sets the iteration marker immediately preceding the suffix.

The functions `fn_composite_name_prepend_comp()` and `fn_composite_name_append_comp()` prepend and append a single component to the given composite name, respectively. These operations invalidate any iteration marker the client holds for that object. `fn_composite_name_insert_comp()` inserts a single component before `iter_pos` to the given composite name and sets `iter_pos` to be immediately after the component just inserted. `fn_composite_name_delete_comp()` deletes the component located before `iter_pos` from the given composite name and sets `iter_pos` back one component.

The functions `fn_composite_name_prepend_name()`, `fn_composite_name_append_name()`, and `fn_composite_name_insert_name()` perform the same update functions as their `_comp` counterparts, respectively, except that multiple components are being added, rather than single components. For example, `fn_composite_name_insert_name()` sets `iter_pos` to be immediately after the name just added.

**RETURN VALUES**

The functions `fn_composite_name_is_empty()`, `fn_composite_name_is_equal()`, `fn_composite_name_is_suffix()`, and `fn_composite_name_is_prefix()` return 1 if the test indicated is true; 0 otherwise.

The update functions `fn_composite_name_prepend_comp()`, `fn_composite_name_append_comp()`, `fn_composite_name_insert_comp()`, `fn_composite_name_delete_comp()`, and their `_name` counterparts return 1 if the update was successful; 0 otherwise.

If a function is expected to return a pointer to an object, a NULL pointer (0) is returned if the function fails.

**ERRORS**

Code set mismatches that occur during the composition of the string form or during comparisons of composite names are resolved in an implementation-dependent way. `fn_string_from_composite_name()`, `fn_composite_name_is_equal()`, `fn_composite_name_is_suffix()`, and `fn_composite_name_is_prefix()` set `status` to `FN_E_INCOMPATIBLE_CODE_SETS` for composite names whose components have code sets that are determined by the implementation to be incompatible.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

## FN\_composite\_name\_t(3XFN)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `FN_string_t(3XFN)`, `xfn(3XFN)`, `attributes(5)`

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

**NAME** FN\_compound\_name\_t, fn\_compound\_name\_from\_syntax\_attrs, fn\_compound\_name\_get\_syntax\_attrs, fn\_compound\_name\_destroy, fn\_string\_from\_compound\_name, fn\_compound\_name\_copy, fn\_compound\_name\_assign, fn\_compound\_name\_count, fn\_compound\_name\_first, fn\_compound\_name\_next, fn\_compound\_name\_prev, fn\_compound\_name\_last, fn\_compound\_name\_prefix, fn\_compound\_name\_suffix, fn\_compound\_name\_is\_empty, fn\_compound\_name\_is\_equal, fn\_compound\_name\_is\_prefix, fn\_compound\_name\_is\_suffix, fn\_compound\_name\_prepend\_comp, fn\_compound\_name\_append\_comp, fn\_compound\_name\_insert\_comp, fn\_compound\_name\_delete\_comp, fn\_compound\_name\_delete\_all – an XFN compound name

**SYNOPSIS**

```
cc [ flag ... ] file ... -lxfn [ library ... ]
#include <xfn/xfn.h>

FN_compound_name_t *fn_compound_name_from_syntax_attrs(const
    FN_attrset_t *aset, const FN_string_t *name, FN_status_t
    *status);

FN_attrset_t *fn_compound_name_get_syntax_attrs(const
    FN_compound_name_t *name);

void fn_compound_name_destroy(FN_compound_name_t *name);

FN_string_t *fn_string_from_compound_name(const
    FN_compound_name_t *name);

FN_compound_name_t *fn_compound_name_copy(const
    FN_compound_name_t *name);

FN_compound_name_t *fn_compound_name_assign(FN_compound_name_t
    *dst, const FN_compound_name_t *src);

unsigned int fn_compound_name_count(const FN_compound_name_t
    *name);

const FN_string_t *fn_compound_name_first(const
    FN_compound_name_t *name, void **iter_pos);

const FN_string_t *fn_compound_name_next(const
    FN_compound_name_t *name, void **iter_pos);

const FN_string_t *fn_compound_name_prev(const
    FN_compound_name_t *name, void **iter_pos);

const FN_string_t *fn_compound_name_last(const
    FN_compound_name_t *name, void **iter_pos);

FN_compound_name_t *fn_compound_name_prefix(const
    FN_compound_name_t *name, const void *iter_pos);

FN_compound_name_t *fn_compound_name_suffix(const
    FN_compound_name_t *name, const void *iter_pos);

int fn_compound_name_is_empty(const FN_compound_name_t *name);
```

## FN\_compound\_name\_t(3XFN)

```
int fn_compound_name_is_equal(const FN_compound_name_t *name1,
    const FN_compound_name_t *name2, unsigned int *status);

int fn_compound_name_is_prefix(const FN_compound_name_t *name,
    const FN_compound_name_t *pre, void **iter_pos, unsigned int
    *status);

int fn_compound_name_is_suffix(const FN_compound_name_t *name,
    const FN_compound_name_t *suffix, void **iter_pos, unsigned int
    *status);

int fn_compound_name_prepend_comp(FN_compound_name_t *name, const
    FN_string_t *atomic_comp, unsigned int *status);

int fn_compound_name_append_comp(FN_compound_name_t *name, const
    FN_string_t *atomic_comp, unsigned int *status);

int fn_compound_name_insert_comp(FN_compound_name_t *name, void
    **iter_pos, const FN_string_t *atomic_comp, unsigned int *status);

int fn_compound_name_delete_comp(FN_compound_name_t *name, void
    **iter_pos);

int fn_compound_name_delete_all(FN_compound_name_t *name);
```

### DESCRIPTION

Most applications treat names as opaque data. Hence, the majority of clients of the XFN interface will not need to parse names. Some applications, however, such as browsers, need to parse names. For these applications, XFN provides support in the form of the `FN_compound_name_t` object.

Each naming system in an XFN federation potentially has its own naming conventions. The `FN_compound_name_t` object has associated operations for applications to process compound names that conform to the XFN model of expressing compound name syntax. The XFN syntax model for compound names covers a large number of specific name syntaxes and is expressed in terms of syntax properties of the naming convention. See `xfn_compound_names(3XFN)`.

An `FN_compound_name_t` object is constructed by the operation `fn_compound_name_from_syntax_attrs`, using a string name and an attribute set containing the "fn\_syntax\_type" (with identifier format `FN_ID_STRING`) attribute identifying the namespace syntax of the string name. The value "standard" (with identifier format `FN_ID_STRING`) in the "fn\_syntax\_type" specifies a syntax model that is by default supported by the `FN_compound_name_t` object. An implementation may support other syntax types instead of the XFN standard syntax model, in which case the value of the "fn\_syntax\_type" attribute would be set to an implementation-specific string. `fn_compound_name_get_syntax_attrs()` returns an attribute set containing the syntax attributes that describes the given compound name. `fn_compound_name_destroy()` releases the storage associated with the given compound name. `fn_string_from_compound_name()` returns the string form of the given compound name. `fn_compound_name_copy()` returns a copy of the given compound name. `fn_compound_name_assign()` makes a copy of the

compound name *src* and assigns it to *dst*, releasing any old contents of *dst*. A pointer to the object pointed to by *dst* is returned. `fn_compound_name_count()` returns the number of atomic components in the given compound name.

The function `fn_compound_name_first()` returns a handle to the `FN_string_t` that is the first atomic component in the compound name, and sets *iter\_pos* to indicate the position immediately following the first component. It returns 0 if the name has no components. Thereafter, successive calls of the `fn_compound_name_next()` function return pointers to the component following the iteration marker, and advance the iteration marker. If the iteration marker is at the end of the sequence, `fn_compound_name_next()` returns 0. Similarly, `fn_compound_name_prev()` returns the component preceding the iteration pointer and moves the marker back one component. If the marker is already at the beginning of the sequence, `fn_compound_name_prev()` returns 0. The function `fn_compound_name_last()` returns a pointer to the last component of the name and sets the iteration marker immediately preceding this component (so that subsequent calls to `fn_compound_name_prev()` can be used to step through trailing components of the name).

The `fn_compound_name_suffix()` function returns a compound name consisting of a copy of those components following the supplied iteration marker. The function `fn_compound_name_prefix()` returns a compound name consisting of those components that precede the iteration marker. Using these functions with an iteration marker that was not initialized with the use of `fn_compound_name_first()`, `fn_compound_name_last()`, `fn_compound_name_is_prefix()`, or `fn_compound_name_is_suffix()` yields undefined and generally undesirable behavior.

The functions `fn_compound_name_is_equal()`, `fn_compound_name_is_prefix()`, and `fn_compound_name_is_suffix()` test for equality between compound names or between parts of compound names. For these functions, equality is defined as name equivalence. A name's syntactic property, such as case-insensitivity, is taken into account by these functions.

The function `fn_compound_name_is_prefix()` tests if one compound name is a prefix of another. If so, it returns 1 and sets the iteration marker immediately following the prefix. (For example, a subsequent call to `fn_compound_name_suffix()` will return the remainder of the name.) Otherwise, it returns 0 and value of the iteration marker is undefined. The function `fn_compound_name_is_suffix()` is similar. It tests if one compound name is a suffix of another. If so, it returns 1 and sets the iteration marker immediately preceding the suffix.

The functions `fn_compound_name_prepend_comp()` and `fn_compound_name_append_comp()` prepend and append a single atomic component to the given compound name, respectively. These operations invalidate any iteration marker the client holds for that object. `fn_compound_name_insert_comp()` inserts an atomic component before *iter\_pos* to the given compound name and sets *iter\_pos* to be immediately after the component

## FN\_compound\_name\_t(3XFN)

just inserted. `fn_compound_name_delete_comp()` deletes the atomic component located before `iter_pos` from the given compound name and sets `iter_pos` back one component. `fn_compound_name_delete_all()` deletes all the atomic components from `name`.

**RETURN VALUES** The following test functions return 1 if the test indicated is true; otherwise, they return 0:

```
fn_compound_name_is_empty()  
fn_compound_name_is_equal()  
fn_compound_name_is_suffix()  
fn_compound_name_is_prefix()
```

The following update functions return 1 if the update was successful; otherwise, they return 0:

```
fn_compound_name_prepend_comp()  
fn_compound_name_append_comp()  
fn_compound_name_insert_comp()  
fn_compound_name_delete_comp()  
fn_compound_name_delete_all()
```

If a function is expected to return a pointer to an object, a NULL pointer (0) is returned if the function fails.

**ERRORS** When the function `fn_compound_name_from_syntax_attrs()` fails, it returns a status code in `status`. The possible status codes are:

<code>FN_E_ILLEGAL_NAME</code>	The name supplied to the operation was not a well- formed XFN compound name, or one of the component names was not well-formed according to the syntax of the naming system(s) involved in its resolution.
<code>FN_E_INCOMPATIBLE_CODE_SETS</code>	The code set of the given string is incompatible with that supported by the compound name.
<code>FN_E_INVALID_SYNTAX_ATTRS</code>	The syntax attributes supplied are invalid or insufficient to fully specify the syntax.
<code>FN_E_SYNTAX_NOT_SUPPORTED</code>	The syntax type specified is not supported.

The following functions may return in `status` the status code `FN_E_INCOMPATIBLE_CODE_SETS` when the code set of the given string is incompatible with that of the compound name:

```
fn_compound_name_is_equal()  
fn_compound_name_is_suffix()  
fn_compound_name_is_prefix()  
fn_compound_name_prepend_comp()  
fn_compound_name_append_comp()
```

FN\_compound\_name\_t(3XFN)

fn\_compound\_name\_insert\_comp()

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_attribute\_t(3XFN), FN\_attrset\_t(3XFN), FN\_composite\_name\_t(3XFN), FN\_status\_t(3XFN), FN\_string\_t(3XFN), fn\_ctx\_get\_syntax\_attrs (3XFN), xfn(3XFN), xfn\_compound\_names(3XFN), attributes(5)

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_ctx\_bind(3XFN)

<b>NAME</b>	fn_ctx_bind – bind a reference to a name				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  int <b>fn_ctx_bind</b>(FN_ctx_t *ctx, const FN_composite_name_t *name,                 const FN_ref_t *ref, unsigned int exclusive, FN_status_t *status);</pre>				
<b>DESCRIPTION</b>	<p>This operation binds the supplied reference <i>ref</i> to the supplied composite name <i>name</i> relative to <i>ctx</i>. The binding is made in the target context, that is, the context named by all but the terminal atomic part of <i>name</i>. The operation binds the terminal atomic name to the supplied reference in the target context. The target context must already exist.</p> <p>The value of <i>exclusive</i> determines what happens if the terminal atomic part of the name is already bound in the target context. If <i>exclusive</i> is nonzero and <i>name</i> is already bound, the operation fails. If <i>exclusive</i> is 0, the new binding replaces any existing binding.</p>				
<b>RETURN VALUES</b>	When the bind operation is successful it returns 1; on error it returns 0.				
<b>ERRORS</b>	fn_ctx_bind sets <i>status</i> as described in FN_status_t(3XFN) and xfn_status_codes. Of special relevance for this operation is the status code FN_E_NAME_IN_USE, which indicates that the supplied name is already in use.				
<b>USAGE</b>	<p>The value of <i>ref</i> cannot be NULL. If the intent is to reserve a name using fn_ctx_bind(), a reference containing no address should be supplied. This reference may be name service-specific or it may be the conventional NULL reference defined in the X/Open registry (see fns_references(5)).</p> <p>If multiple sources are updating a reference, they must synchronize amongst each other when adding, modifying, or removing from the address list of a bound reference.</p>				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	FN_composite_name_t(3XFN), FN_ctx_t(3XFN), FN_ref_t(3XFN), FN_status_t(3XFN), fn_ctx_lookup(3XFN), fn_ctx_unbind(3XFN), xfn(3XFN), xfn_status_codes(3XFN), attributes(5), fns_references(5)				
<b>NOTES</b>	The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As				



fn\_ctx\_bind(3XFN)

the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_ctx\_create\_subcontext(3XFN)

<b>NAME</b>	fn_ctx_create_subcontext – create a subcontext in a context				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_ref_t *fn_ctx_create_subcontext(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);</pre>				
<b>DESCRIPTION</b>	<p>This operation creates a new XFN context of the same type as the target context — that named by all but the terminal atomic component of <i>name</i> — and binds it to the supplied composite name.</p> <p>As with <code>fn_ctx_bind( )</code>, the target context must already exist. The new context is created and bound in the target context using the terminal atomic name in <i>name</i>. The operation returns a reference to the newly created context.</p>				
<b>RETURN VALUE</b>	fn_ctx_create_subcontext( ) returns a reference to the newly created context; if the operation fails, it returns a NULL pointer (0).				
<b>ERRORS</b>	<p>fn_ctx_create_subcontext( ) sets <i>status</i> as described in FN_status_t(3XFN) and xfn_status_codes(3XFN). Of special relevance for this operation is the following status code:</p> <table border="0"> <tr> <td style="padding-right: 20px;">FN_E_NAME_IN_USE</td> <td>The terminal atomic name already exists in the target context.</td> </tr> </table>	FN_E_NAME_IN_USE	The terminal atomic name already exists in the target context.		
FN_E_NAME_IN_USE	The terminal atomic name already exists in the target context.				
<b>APPLICATION USAGE</b>	The new subcontext is an XFN context and is created in the same naming system as the target context. The new subcontext also inherits the same syntax attributes as the target context. XFN does not specify any further properties of the new subcontext. The target context and its naming system determine these.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe.</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe.
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe.				
<b>SEE ALSO</b>	FN_composite_name_t(3XFN), FN_ctx_t(3XFN), FN_ref_t(3XFN), FN_status_t(3XFN), fn_ctx_bind(3XFN), fn_ctx_lookup(3XFN), fn_ctx_destroy_subcontext(3XFN), xfn_status_codes(3XFN), xfn(3XFN), attributes(5)				

fn\_ctx\_destroy\_subcontext(3XFN)

**NAME** | fn\_ctx\_destroy\_subcontext – destroy the named context and remove its binding from the parent context

**SYNOPSIS** | 

```
cc [ flag ... ] file ... -lxfn [ library ... ]
#include <xfn/xfn.h>

int fn_ctx_destroy_subcontext(FN_ctx_t *ctx, const
    FN_composite_name_t *name, FN_status_t *status);
```

**DESCRIPTION** | This operation destroys the subcontext named by *name* relative to *ctx*, and unbinds the name.

As with `fn_ctx_unbind( )`, this operation succeeds even if the terminal atomic name is not bound in the target context — the context named by all but the terminal atomic name in *name*.

**RETURN VALUE** | `fn_ctx_destroy_subcontext( )` returns 1 on success and 0 on failure.

**ERRORS** | `fn_ctx_destroy_subcontext( )` sets *status* as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`. Of special relevance for `fn_ctx_destroy_subcontext( )` are the following status codes:

`FN_E_CTX_NOT_A_CONTEXT` *name* does not name a context.

`FN_E_CTX_NOT_EMPTY` The naming system being asked to do the destroy does not support removal of a context that still contains bindings.

**APPLICATION USAGE** | Some aspects of this operation are not specified by XFN, but are determined by the target context and its naming system. For example, XFN does not specify what happens if the named subcontext is non-empty when the operation is invoked.

In naming systems that support attributes, and store the attributes along with names or contexts, this operation removes the name, the context, and its associated attributes.

Normal resolution always follows links. In a `fn_ctx_destroy_subcontext( )` operation, resolution of *name* continues to the target context; the terminal atomic name is not resolved. If the terminal atomic name is bound to a link, the link is not followed and the operation fails with `FN_E_CTX_NOT_A_CONTEXT` because the name is not bound to a context.

**ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe.

**SEE ALSO** | `FN_ctx_t(3XFN)`, `FN_composite_name_t(3XFN)`, `FN_status_t(3XFN)`, `fn_ctx_create_subcontext(3XFN)`, `fn_ctx_unbind(3XFN)`, `xfn(3XFN)`, `xfn_status_codes(3XFN)`, `attributes(5)`

## fn\_ctx\_equivalent\_name(3XFN)

<b>NAME</b>	fn_ctx_equivalent_name – construct an equivalent name in same context		
<b>SYNOPSIS</b>	<pre>#include &lt;xfn/xfn.h&gt;  FN_composite_name_t *fn_ctx_equivalent_name(FN_ctx_t *ctx, const       FN_composite_name_t *name, const FN_string_t *leading_name,       FN_status_t *status);</pre>		
<b>DESCRIPTION</b>	<p>Given the name of an object <i>name</i> relative to the context <i>ctx</i>, this operation returns an equivalent name for that object, relative to the same context <i>ctx</i>, that has <i>leading_name</i> as its initial atomic name. Two names are said to be equivalent if they have prefixes that resolve to the same context, and the parts of the names immediately following the prefixes are identical.</p> <p>The existence of a binding for <i>leading_name</i> in <i>ctx</i> does not guarantee that a name equivalent to <i>name</i> can be constructed. The failure may be because such equivalence is not meaningful, or due to the inability of the system to construct a name with the equivalence. For example, supplying <code>_thishost</code> as <i>leading_name</i> when <i>name</i> starts with <code>_myself</code> to <code>fn_ctx_equivalent_name()</code> in the Initial Context would not be meaningful; this results in the return of the error code <code>FN_E_NO_EQUIVALENT_NAME</code>.</p>		
<b>RETURN VALUES</b>	If an equivalent name cannot be constructed, the value 0 is returned and <i>status</i> is set appropriately.		
<b>ERRORS</b>	<p><code>fn_ctx_equivalent_name()</code> sets <i>status</i> as described in <code>FN_status_t(3XFN)</code> and <code>xfn_status_codes(3XFN)</code>. The following status code is especially relevant for this operation:</p> <table><tr><td><code>FN_E_NO_EQUIVALENT_NAME</code></td><td>No equivalent name can be constructed, either because there is no meaningful equivalence between <i>name</i> and <i>leading_name</i>, or the system does not support constructing the requested equivalent name, for implementation-specific reasons.</td></tr></table>	<code>FN_E_NO_EQUIVALENT_NAME</code>	No equivalent name can be constructed, either because there is no meaningful equivalence between <i>name</i> and <i>leading_name</i> , or the system does not support constructing the requested equivalent name, for implementation-specific reasons.
<code>FN_E_NO_EQUIVALENT_NAME</code>	No equivalent name can be constructed, either because there is no meaningful equivalence between <i>name</i> and <i>leading_name</i> , or the system does not support constructing the requested equivalent name, for implementation-specific reasons.		
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Naming Files</p> <p>In the Initial Context supporting XFN enterprise policies, a user <code>jsmith</code> is able to name one of her files relative to this context in several ways.</p> <pre>_myself/_fs/map.ps _user/jsmith/_fs/map.ps _orgunit/finance/_user/jsmith/_fs/map.ps</pre> <p>The first of these may be appealing to the user <code>jsmith</code> in her day-to-day operations. This name is not, however, appropriate for her to use when referring the file in an electronic mail message sent to a colleague. The second of these names would be appropriate if the colleague were in the same organizational unit, and the third appropriate for anyone in the same enterprise.</p>		

**EXAMPLE 1** Naming Files *(Continued)*

When the following sequence of instructions is executed by the user `jsmith` in the organizational unit `finance`, `enterprise_wide_name` would contain the composite name `_orgunit/finance/_user/jsmith/_fs/map.ps`:

```
FN_string_t* namestr =
    fn_string_from_str((const unsigned char*)_myself/_fs/map.ps");
FN_composite_name_t* name = fn_composite_name_from_string(namestr);
FN_string_t* org_lead =
    fn_string_from_str((const unsigned char*)_orgunit");
FN_status_t* status = fn_status_create();
FN_composite_name_t* enterprise_wide_name;
FN_ctx_t* init_ctx = fn_ctx_handle_from_initial(status);
/* check status of from_initial( ) */
enterprise_wide_name = fn_ctx_equivalent_name(init_ctx, name, org_lead,
status);
```

When the following sequence of instructions is executed by the user `jsmith` in the organizational unit `finance`, `shortest_name` would contain the composite name `_myself/_fs/map.ps`:

```
FN_string_t* namestr =
    fn_string_from_str((const unsigned char*)
        "_orgunit/finance/_user/jsmith/_fs/map.ps");
FN_composite_name_t* name = fn_composite_name_from_string(namestr);
FN_string_t* mylead = fn_string_from_str((const unsigned char*)_myself");
FN_status_t* status = fn_status_create();
FN_composite_name_t* shortest_name;
FN_ctx_t* init_ctx = fn_ctx_handle_from_initial(status);
/* check status of from_initial( ) */
shortest_name = fn_ctx_equivalent_name(init_ctx, name, mylead, status);
```

**ATTRIBUTES** See attributes (5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `FN_composite_name_t(3XFN)`, `FN_ctx_t(3XFN)`, `FN_status_t(3XFN)`, `FN_string_t(3XFN)`, `xfn_status_codes(3XFN)`, `attributes(5)`

## fn\_ctx\_get\_ref(3XFN)

<b>NAME</b>	fn_ctx_get_ref – return a context’s reference				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_ref_t *fn_ctx_get_ref(const FN_ctx_t *ctx, FN_status_t *status);</pre>				
<b>DESCRIPTION</b>	This operation returns a reference to the supplied context object.				
<b>RETURN VALUE</b>	fn_ctx_get_ref() returns a pointer to an FN_ref_t object if the operation succeeds, it returns 0 if the operation fails.				
<b>ERRORS</b>	fn_ctx_get_ref() sets <i>status</i> as described in FN_status_t(3XFN) and xfn_status_codes(3XFN). The following status code is of particular relevance to this operation:  FN_E_OPERATION_NOT_SUPPORTED      Using the fn_ctx_get_ref() operation on the Initial Context returns this status code.				
<b>APPLICATION USAGE</b>	fn_ctx_get_ref() cannot be used on the Initial Context. fn_ctx_get_ref() can be used on contexts bound in the Initial Context (in other words, the bindings in the Initial Context have references).  If the context handle was created earlier using the fn_ctx_handle_from_ref() operation, the reference returned by the fn_ctx_get_ref() operation may not necessarily be exactly the same in content as that originally supplied. For example, fn_ctx_handle_from_ref() may construct the context handle from one address from the list of addresses. The context implementation may return with a call to fn_ctx_get_ref() only that address, or a more complete list of addresses than what was supplied in fn_ctx_handle_from_ref().				
<b>ATTRIBUTES</b>	See attributes (5) for descriptions of the following attributes:  <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>Safe.</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe.
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe.				
<b>SEE ALSO</b>	FN_ctx_t(3XFN), FN_ref_t(3XFN), FN_status_t(3XFN), fn_ctx_handle_from_initial(3XFN), fn_ctx_handle_from_ref(3XFN), xfn_status_codes (3XFN), xfn(3XFN), attributes(5)				

<b>NAME</b>	fn_ctx_get_syntax_attrs – return syntax attributes associated with named context
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_attrset_t *fn_ctx_get_syntax_attrs(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);</pre>
<b>DESCRIPTION</b>	<p>Each context has an associated set of syntax-related attributes. This operation returns the syntax attributes associated with the context named by <i>name</i> relative to the context <i>ctx</i>.</p> <p>The attributes must contain the attribute <code>fn_syntax_type</code> ( <code>FN_ID_STRING</code> format). If the context supports a syntax that conforms to the XFN standard syntax model, <code>fn_syntax_type</code> is set to "standard" (ASCII attribute syntax) and the attribute set contains the rest of the relevant syntax attributes described in <code>xfn_compound_names(3XFN)</code>.</p> <p>This operation is different from other XFN attribute operations in that these syntax attributes could be obtained directly from the context. Attributes obtained through other XFN attribute operations may not necessarily be associated with the context; they may be associated with the reference of context, rather than the context itself (see <code>xfn_attributes(3XFN)</code>).</p>
<b>RETURN VALUE</b>	<code>fn_ctx_get_syntax_attrs()</code> returns an attribute set if successful; it returns a NULL pointer (0) if the operation fails.
<b>ERRORS</b>	<code>fn_ctx_get_syntax_attrs()</code> sets <i>status</i> as described in <code>FN_status_t(3XFN)</code> and <code>xfn_status_codes(3XFN)</code> .
<b>APPLICATION USAGE</b>	<p>Implementations may choose to support other syntax types in addition to, or in place of, the XFN standard syntax model, in which case, the value of the <code>fn_syntax_type</code> attribute would be set to an implementation-specific string, and different or additional syntax attributes will be in the set.</p> <p>Syntax attributes of a context may be generated automatically by a context, in response to <code>fn_ctx_get_syntax_attrs()</code>, or they may be created and updated using the base attribute operations. This is implementation-dependent.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe.

fn\_ctx\_get\_syntax\_attrs(3XFN)

**SEE ALSO** | FN\_attrset\_t(3XFN), FN\_composite\_name\_t(3XFN),  
FN\_compound\_name\_t(3XFN), FN\_ctx\_t(3XFN), FN\_status\_t(3XFN),  
fn\_attr\_get(3XFN), fn\_attr\_multi\_get(3XFN),  
xfn\_compound\_names(3XFN), xfn\_attributes(3XFN),  
xfn\_status\_codes(3XFN), xfn(3XFN), attributes(5)



fn\_ctx\_handle\_destroy(3XFN)

**NAME** | fn\_ctx\_handle\_destroy – release storage associated with context handle

**SYNOPSIS** | 

```
cc [ flag ... ] file ... -lxfn [ library ... ]
#include <xfn/xfn.h>

void fn_ctx_handle_destroy(FN_ctx_t *ctx);
```

**DESCRIPTION** | This operation destroys the context handle *ctx* and allows the implementation to free resources associated with the context handle. This operation does not affect the state of the context itself.

**ATTRIBUTES** | See attributes (5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe.

**SEE ALSO** | FN\_ctx\_t(3XFN), fn\_ctx\_handle\_from\_initial(3XFN),  
fn\_ctx\_handle\_from\_ref(3XFN), xfn(3XFN), attributes(5)

## fn\_ctx\_handle\_from\_initial(3XFN)

<b>NAME</b>	fn_ctx_handle_from_initial – return a handle to the Initial Context
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_ctx_t *fn_ctx_handle_from_initial(unsigned int authoritative,                                      FN_status_t *status);</pre>
<b>DESCRIPTION</b>	<p>This operation returns a handle to the caller's Initial Context. On successful return, the handle points to a context which meets the specification of the XFN Initial Context (see <code>fns_initial_context(5)</code>).</p> <p><i>authoritative</i> specifies whether the handle to the context returned should be authoritative with respect to information the context obtains from the naming service. When the flag is non-zero, subsequent operations on the context will access the most authoritative information. When <i>authoritative</i> is 0, the handle to the context returned need not be authoritative.</p>
<b>RETURN VALUES</b>	fn_ctx_handle_from_initial() returns a pointer to an FN_ctx_t object if the operation succeeds; it returns a NULL pointer (0) otherwise.
<b>ERRORS</b>	fn_ctx_handle_from_initial() sets only the status code portion of the status object <i>status</i> .
<b>USAGE</b>	<p>Authoritativeness is determined by specific naming services. For example, in a naming service that supports replication using a master/slave model, the source of authoritative information would come from the master server. In some naming systems, bypassing the naming service cache may reach servers which provide the most authoritative information. The availability of an authoritative context might be lower due to the lower number of servers offering this service. For the same reason, it might also provide poorer performance than contexts that need not be authoritative.</p> <p>Applications set <i>authoritative</i> to 0 for typical day-to-day operations. Applications only set <i>authoritative</i> to a non-zero value when they require access to the most authoritative information, possibly at the expense of lower availability and/or poorer performance.</p> <p>It is implementation-dependent whether authoritativeness is transferred from one context to the next as composite name resolution proceeds. Getting an authoritative context handle to the Initial Context means that operations on bindings in the Initial Context are processed using the most authoritative information. Contexts referenced implicitly through an authoritative Initial Context (for example, through the use of composite names) may not necessarily themselves be authoritative.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

fn\_ctx\_handle\_from\_initial(3XFN)

**SEE ALSO** FN\_ctx\_t(3XFN), FN\_status\_t(3XFN), fn\_ctx\_get\_ref(3XFN),  
fn\_ctx\_handle\_from\_ref(3XFN), xfn(3XFN), xfn\_status\_codes(3XFN),  
attributes(5), fns\_initial\_context(5)

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_ctx\_handle\_from\_ref(3XFN)

<b>NAME</b>	fn_ctx_handle_from_ref – construct a handle to a context object using the given reference		
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_ctx_t *fn_ctx_handle_from_ref(const FN_ref_t *ref, unsigned int     authoritative, FN_status_t *status);</pre>		
<b>DESCRIPTION</b>	<p>This operation creates a handle to an FN_ctx_t object using an FN_ref_t object for that context.</p> <p><i>authoritative</i> specifies whether the handle to the context returned should be authoritative with respect to information the context obtains from the naming service. When the flag is non-zero, subsequent operations on the context will access the most authoritative information. When <i>authoritative</i> is 0, the handle to the context returned need not be authoritative.</p>		
<b>RETURN VALUES</b>	This operation returns a pointer to an FN_ctx_t object if the operation succeeds; otherwise, it returns a NULL pointer (0).		
<b>ERRORS</b>	<p>fn_ctx_handle_from_ref() sets <i>status</i> as described in FN_status_t(3XFN) and xfn_status_codes(3XFN). The following status code is of particular relevance to this operation:</p> <table><tr><td>FN_E_NO_SUPPORTED_ADDRESS</td><td>A context object could not be constructed from a particular reference. The reference contained no address type over which the context interface was supported.</td></tr></table>	FN_E_NO_SUPPORTED_ADDRESS	A context object could not be constructed from a particular reference. The reference contained no address type over which the context interface was supported.
FN_E_NO_SUPPORTED_ADDRESS	A context object could not be constructed from a particular reference. The reference contained no address type over which the context interface was supported.		
<b>USAGE</b>	<p>Authoritativeness is determined by specific naming services. For example, in a naming service that supports replication using a master/slave model, the source of authoritative information would come from the master server. In some naming systems, bypassing the naming service cache may reach servers which provide the most authoritative information. The availability of an authoritative context might be lower due to the lower number of servers offering this service. For the same reason, it might also provide poorer performance than contexts that need not be authoritative.</p> <p>Applications set <i>authoritative</i> to 0 for typical day-to-day operations. Applications only set <i>authoritative</i> to a non-zero value when they require access to the most authoritative information, possibly at the expense of lower availability and/or poorer performance.</p> <p>To control the authoritativeness of the target context, the application first resolves explicitly to the target context using fn_ctx_lookup(3XFN). It then uses fn_ctx_handle_from_ref() with the appropriate authoritative argument to obtain a handle to the context. This returns a handle to a context with the specified authoritativeness. The application then uses the XFN operations, such as lookup and list, with this context handle.</p>		

fn\_ctx\_handle\_from\_ref(3XFN)

It is implementation-dependent whether authoritativeness is transferred from one context to the next as composite name resolution proceeds. The application should use the approach recommended above to achieve the desired level of authoritativeness on a per context basis.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `FN_ctx_t(3XFN)`, `FN_ref_t(3XFN)`, `FN_status_t(3XFN)`, `fn_ctx_get_ref(3XFN)`, `fn_ctx_handle_destroy(3XFN)`, `fn_ctx_lookup(3XFN)`, `xfn(3XFN)`, `xfn_status_codes(3XFN)`, `attributes(5)`, `fns_references(5)`

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_ctx\_list\_bindings(3XFN)

<b>NAME</b>	fn_ctx_list_bindings, FN_bindinglist_t, fn_bindinglist_next, fn_bindinglist_destroy – list the atomic names and references bound in a context				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_bindinglist_t *fn_ctx_list_bindings(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);  FN_string_t *fn_bindinglist_next(FN_bindinglist_t *bl, FN_ref_t     **ref, FN_status_t *status);  void fn_bindinglist_destroy(FN_bindinglist_t *bl, FN_status_t     *status);</pre>				
<b>DESCRIPTION</b>	<p>This set of operations is used to list the names and bindings in the context named by <i>name</i> relative to the context <i>ctx</i>. Note that <i>name</i> must name a context. If the intent is to list the contents of <i>ctx</i>, <i>name</i> should be an empty composite name.</p> <p>The semantics of these operations are similar to those for listing names (see <code>fn_ctx_list_names(3XFN)</code>). In addition to a name string being returned, <code>fn_bindinglist_next()</code> also returns the reference of the binding for each member of the enumeration.</p>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>FN_composite_name_t(3XFN)</code> , <code>FN_ctx_t(3XFN)</code> , <code>FN_ref_t(3XFN)</code> , <code>FN_status_t(3XFN)</code> , <code>FN_string_t(3XFN)</code> , <code>fn_ctx_list_names(3XFN)</code> , <code>xfn(3XFN)</code> , <code>xfn_status_codes(3XFN)</code> , <code>attributes(5)</code>				
<b>NOTES</b>	The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.				

<b>NAME</b>	fn_ctx_list_names, FN_namelist_t, fn_namelist_next, fn_namelist_destroy – list the atomic names bound in a context				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_namelist_t *fn_ctx_list_names(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);  FN_string_t *fn_namelist_next(FN_namelist_t *nl, FN_status_t     *status);  void fn_namelist_destroy(FN_namelist_t *nl, FN_status_t *status);</pre>				
<b>DESCRIPTION</b>	<p>This set of operations is used to list the names bound in the target context named <i>name</i> relative to the context <i>ctx</i>. Note that <i>name</i> must name a context. If the intent is to list the contents of <i>ctx</i>, <i>name</i> should be an empty composite name.</p> <p>The call to <code>fn_ctx_list_names()</code> initiates the enumeration process. It returns a handle to an <code>FN_namelist_t</code> object that can be used to enumerate the names in the target context.</p> <p>The operation <code>fn_namelist_next()</code> returns the next name in the enumeration identified by <i>nl</i> and updates <i>nl</i> to indicate the state of the enumeration. Successive calls to <code>fn_namelist_next()</code> using <i>nl</i> return successive names in the enumeration and further update the state of the enumeration. <code>fn_namelist_next()</code> returns a NULL pointer (0) when the enumeration has been completed.</p> <p><code>fn_namelist_destroy()</code> is used to release resources used during the enumeration. This may be invoked at any time to terminate the enumeration.</p>				
<b>RETURN VALUES</b>	<p><code>fn_ctx_list_names()</code> returns a pointer to an <code>FN_namelist_t</code> object if the enumeration is successfully initiated; otherwise it returns a NULL pointer (0).</p> <p><code>fn_namelist_next()</code> returns a NULL pointer (0) if no more names can be returned in the enumeration.</p> <p>In the case of a failure, these operations return in <i>status</i> a code indicating the nature of the failure.</p>				
<b>ERRORS</b>	<p>Each successful call to <code>fn_namelist_next()</code> returns a name and sets <i>status</i> to <code>FN_SUCCESS</code>.</p> <p>When <code>fn_namelist_next()</code> returns a NULL pointer (0), it indicates that no more names can be returned. <i>status</i> is set in the following way:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;"><code>FN_SUCCESS</code></td> <td>The enumeration has completed successfully.</td> </tr> <tr> <td><code>FN_E_INVALID_ENUM_HANDLE</code></td> <td>The supplied enumeration handle is not valid. Possible reasons could be that the handle was from another enumeration, or the context being enumerated no longer</td> </tr> </table>	<code>FN_SUCCESS</code>	The enumeration has completed successfully.	<code>FN_E_INVALID_ENUM_HANDLE</code>	The supplied enumeration handle is not valid. Possible reasons could be that the handle was from another enumeration, or the context being enumerated no longer
<code>FN_SUCCESS</code>	The enumeration has completed successfully.				
<code>FN_E_INVALID_ENUM_HANDLE</code>	The supplied enumeration handle is not valid. Possible reasons could be that the handle was from another enumeration, or the context being enumerated no longer				

## fn\_ctx\_list\_names(3XFN)

accepts the handle (due to such events as handle expiration or updates to the context).

FN\_E\_PARTIAL\_RESULT

The enumeration is not yet complete but cannot be continued.

Other status codes, such as FN\_E\_COMMUNICATION\_FAILURE, are also possible in calls to `fn_ctx_list_names()`, `fn_namelist_next()`, and `fn_namelist_destroy()`. These functions set *status* for these other status codes as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`.

### USAGE

The names enumerated using `fn_namelist_next()` are not ordered in any way. There is no guaranteed relation between the order in which names are added to a context and the order of names obtained by enumeration. The specification does *not* guarantee that any two series of enumerations will return the names in the same order.

When a name is added to or removed from a context, this may or may not invalidate the enumeration handle that the client holds for that context. If the enumeration handle becomes invalid, the status code FN\_E\_INVALID\_ENUM\_HANDLE is returned in *status*. If the enumeration handle remains valid, the update may or may not be visible to the client.

In addition, there may be a relationship between the *ctx* argument supplied to `fn_ctx_list_names()` and the `FN_namelist_t` object it returns. For example, some implementations may store the context handle *ctx* within the `FN_namelist_t` object for subsequent `fn_namelist_next()` calls. In general, a `fn_ctx_handle_destroy(3XFN)` should not be invoked on *ctx* until the enumeration has terminated.

### EXAMPLES

**EXAMPLE 1** A sample program.

The following code fragment illustrates how the list names operations may be used:

```
extern FN_string_t *user_input;
FN_ctx_t *ctx;
FN_composite_name_t *target_name = fn_composite_name_from_string(user_input);
FN_status_t *status = fn_status_create();
FN_string_t *name;
FN_namelist_t *nl;
ctx = fn_ctx_handle_from_initial(status);
/* error checking on 'status' */
if ((nl=fn_ctx_list_names(ctx, target_name, status)) == 0) {
    /* report 'status' and exit */
}
while (name=fn_namelist_next(nl, status)) {
    /* do something with 'name' */
    fn_string_destroy(name);
}
/* check 'status' for reason for end of enumeration and report if necessary */
/* clean up */
```



**EXAMPLE 1** A sample program. (Continued)

```
fn_namelist_destroy(nl, status);
/* report 'status' */
```

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_composite\_name\_t(3XFN), FN\_ctx\_t(3XFN), FN\_status\_t(3XFN), FN\_string\_t(3XFN), fn\_ctx\_handle\_destroy(3XFN), xfn(3XFN), xfn\_status\_codes(3XFN), attributes(5)

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_ctx\_lookup(3XFN)

**NAME** | fn\_ctx\_lookup – look up name in context

**SYNOPSIS** | 

```
cc [ flag ... ] file ... -lxfn [ library ... ]
#include <xfn/xfn.h>

FN_ref_t *fn_ctx_lookup(FN_ctx_t *ctx, const FN_composite_name_t
    *name, FN_status_t *status);
```

**DESCRIPTION** | This operation returns the reference bound to *name* relative to the context *ctx*.

**RETURN VALUE** | If the operation succeeds, the `fn_ctx_lookup()` function returns a handle to the reference bound to *name*. Otherwise, 0 is returned and *status* is set appropriately.

**ERRORS** | `fn_ctx_lookup()` sets *status* as described `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`.

**APPLICATION USAGE** | Some naming services may not always have reference information for all names in their contexts; for such names, such naming services may return a special reference whose type indicates that the name is not bound to any address. This reference may be name service specific or it may be the conventional NULL reference defined in the X/Open registry. See `fns_references(5)`.

**ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe.

**SEE ALSO** | `FN_composite_name_t(3XFN)`, `FN_ctx_t(3XFN)`, `FN_ref_t(3XFN)`, `FN_status_t(3XFN)`, `fns_references(5)`, `xfn_status_codes(3XFN)`, `xfn(3XFN)`, `attributes(5)`

- NAME** fn\_ctx\_lookup\_link – look up the link reference bound to a name
- SYNOPSIS**

```
cc [ flag ... ] file ... -lxfn [ library ... ]
#include <xfn/xfn.h>

FN_ref_t *fn_ctx_lookup_link(FN_ctx_t *ctx, const
    FN_composite_name_t *name, FN_status_t *status);
```
- DESCRIPTION** This operation returns the XFN link bound to *name*. The terminal atomic part of *name* must be bound to an XFN link.
- The normal fn\_ctx\_lookup(3XFN) operation follows all links encountered, including any bound to the terminal atomic part of *name*. This operation differs from the normal lookup in that when the terminal atomic part of *name* is an XFN link, this link is not followed, and the operation returns the link.
- RETURN VALUES** If fn\_ctx\_lookup\_link() fails, a NULL pointer (0) is returned.
- ERRORS** fn\_ctx\_lookup\_link() sets *status* as described in FN\_status\_t(3XFN) and xfn\_status\_codes(3XFN). Of special relevance for fn\_ctx\_lookup\_link() is the following status code:
- FN\_E\_MALFORMED\_LINK      *name* resolved to a reference that was not a link.
- ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_composite\_name\_t(3XFN), FN\_ctx\_t(3XFN), FN\_ref\_t(3XFN), FN\_status\_t(3XFN), fn\_ctx\_lookup(3XFN), xfn(3XFN), xfn\_links(3XFN), xfn\_status\_codes(3XFN), attributes(5)

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## fn\_ctx\_rename(3XFN)

<b>NAME</b>	fn_ctx_rename – rename the name of a binding				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  int <b>fn_ctx_rename</b>(FN_ctx_t *ctx, const FN_composite_name_t     *oldname, const FN_composite_name_t *newname, unsigned int     exclusive, FN_status_t *status);</pre>				
<b>DESCRIPTION</b>	<p>The <code>fn_ctx_rename()</code> operation binds the reference currently bound to <code>oldname</code> relative to <code>ctx</code>, to the name <code>newname</code>, and unbinds <code>oldname</code>. <code>newname</code> is resolved relative to the target context (that named by all but the terminal atomic part of <code>oldname</code>).</p> <p>If <code>exclusive</code> is 0, the operation overwrites any old binding of <code>newname</code>. If <code>exclusive</code> is nonzero, the operation fails if <code>newname</code> is already bound.</p>				
<b>RETURN VALUES</b>	<code>fn_ctx_rename()</code> returns 1 if the operation is successful, 0 otherwise.				
<b>ERRORS</b>	<code>fn_ctx_rename()</code> sets <code>status</code> as described <code>FN_status_t(3XFN)</code> and <code>xfn_status_codes(3XFN)</code> .				
<b>USAGE</b>	<p>The only restriction that XFN places on <code>newname</code> is that it be resolved relative to the target context. XFN does not specify further restrictions on <code>newname</code>. For example, in some implementations, <code>newname</code> might be restricted to be a name in the same naming system as the terminal component of <code>oldname</code>. In another implementation, <code>newname</code> might be restricted to be an atomic name.</p> <p>Normal resolution always follows links. In an <code>fn_ctx_rename()</code> operation, resolution of <code>oldname</code> continues to the target context; the terminal atomic name is not resolved. If the terminal atomic name is bound to a link, the link is not followed and the operation binds <code>newname</code> to the link and unbinds the terminal atomic name of <code>oldname</code>.</p> <p>In naming systems that support attributes and store the attributes along with the names, the unbind of the terminal atomic name of <code>oldname</code> also removes its associated attributes. It is implementation-dependent whether these attributes become associated with <code>newname</code>.</p>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>FN_composite_name_t(3XFN)</code> , <code>FN_ctx_t(3XFN)</code> , <code>FN_ref_t(3XFN)</code> , <code>FN_status_t(3XFN)</code> , <code>fn_ctx_bind(3XFN)</code> , <code>fn_ctx_unbind(3XFN)</code> , <code>xfn(3XFN)</code> , <code>xfn_status_codes(3XFN)</code> , <code>attributes(5)</code>				

fn\_ctx\_rename(3XFN)

**NOTES** | The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## FN\_ctx\_t(3XFN)

<b>NAME</b>	FN_ctx_t – an XFN context
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_ctx_t *fn_ctx_handle_from_initial(unsigned int authoritative,     FN_status_t *status);  FN_ctx_t *fn_ctx_handle_from_ref(const FN_ref_t *ref, unsigned int     authoritative, FN_status_t *status);  FN_ref_t *fn_ctx_get_ref(const FN_ctx_t *ctx, FN_status_t *status); void fn_ctx_handle_destroy(FN_ctx_t *ctx);  FN_ref_t *fn_ctx_lookup(FN_ctx_t *ctx, const FN_composite_name_t     *name, FN_status_t *status);  FN_namelist_t *fn_ctx_list_names(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);  FN_string_t *fn_namelist_next(FN_namelist_t *nl, FN_status_t     *status);  void fn_namelist_destroy(FN_namelist_t *nl, FN_status_t *status);  FN_bindinglist_t *fn_ctx_list_bindings(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);  FN_string_t *fn_bindinglist_next(FN_bindinglist_t *iter, FN_ref_t     **ref, FN_status_t *status);  void fn_bindinglist_destroy(FN_bindinglist_t *iter_pos, FN_status_t     *status);  int fn_ctx_bind(FN_ctx_t *ctx, const FN_composite_name_t *name,     const FN_ref_t *ref, unsigned int exclusive, FN_status_t *status);  int fn_ctx_unbind(FN_ctx_t *ctx, const FN_composite_name_t *name,     FN_status_t *status);  int fn_ctx_rename(FN_ctx_t *ctx, const FN_composite_name_t     *oldname, const FN_composite_name_t *newname, unsigned int     exclusive, FN_status_t *status);  FN_ref_t *fn_ctx_create_subcontext(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);  int fn_ctx_destroy_subcontext(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);  FN_ref_t *fn_ctx_lookup_link(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);  FN_attrset_t *fn_ctx_get_syntax_attrs(FN_ctx_t *ctx, const     FN_composite_name_t *name, FN_status_t *status);</pre>

<b>DESCRIPTION</b>	<p>An XFN context consists of a set of name to reference bindings. An XFN context is represented by the type <code>FN_ctx_t</code> in the client interface. The operations for manipulating an <code>FN_ctx_t</code> object are described in detail in separate reference manual pages.</p> <p>The following contains a brief summary of these operations:</p> <p><code>fn_ctx_handle_from_initial()</code> returns a pointer to an Initial Context that provides a starting point for resolution of composite names.</p> <p><code>fn_ctx_handle_from_ref()</code> returns a handle to an <code>FN_ctx_t</code> object using the given reference <i>ref</i>. <code>fn_ctx_get_ref()</code> returns the reference of the context <i>ctx</i>.</p> <p><code>fn_ctx_handle_destroy()</code> releases the resources associated with the <code>FN_ctx_t</code> object <i>ctx</i>; it does not affect the state of the context itself.</p> <p><code>fn_ctx_lookup()</code> returns the reference bound to <i>name</i> resolved relative to <i>ctx</i>.</p> <p><code>fn_ctx_list_names()</code> is used to enumerate the atomic names bound in the context named by <i>name</i> resolved relative to <i>ctx</i>. <code>fn_ctx_list_bindings()</code> is used to enumerate the atomic names and their references in the context named by <i>name</i> resolved relative to <i>ctx</i>.</p> <p><code>fn_ctx_bind()</code> binds the composite name <i>name</i> to a reference <i>ref</i> resolved relative to <i>ctx</i>. <code>fn_ctx_unbind()</code> unbinds <i>name</i> resolved relative to <i>ctx</i>. <code>fn_ctx_rename()</code> binds <i>newname</i> to the reference bound to <i>oldname</i> and unbinds <i>oldname</i>. <i>oldname</i> is resolved relative to <i>ctx</i>; <i>newname</i> is resolved relative to the target context.</p> <p><code>fn_ctx_create_subcontext()</code> creates a new context with the given composite name <i>name</i> resolved relative to <i>ctx</i>. <code>fn_ctx_destroy_subcontext()</code> destroys the context named by <i>name</i> resolved relative to <i>ctx</i>.</p> <p>Normal resolution always follows links. <code>fn_ctx_lookup_link()</code> looks up <i>name</i> relative to <i>ctx</i>, following links except for the last atomic part of <i>name</i>, which must be bound to an XFN link.</p> <p><code>fn_ctx_get_syntax_attrs()</code> returns an attribute set containing attributes that describe a context's syntax. <i>name</i> must name a context.</p>
<b>ERRORS</b>	<p>In each context operation, the caller supplies an <code>FN_status_t</code> object as a parameter. The called function sets this status object as described in <code>FN_status_t(3XFN)</code> and <code>xfn_status_codes(3XFN)</code>.</p>
<b>USAGE</b>	<p>In most of the operations of the base context interface, the caller supplies a context and a composite name. The supplied name is always interpreted relative to the supplied context.</p> <p>The operation may eventually be effected on a different context called the operation's <i>target context</i>. Each operation has an initial resolution phase that conveys the operation to its target context, and the operation is then applied. The effect (but not necessarily the implementation) is that of doing a lookup on that portion of the name that represents the target context, and then invoking the operation on the target context. The contexts involved only in the resolution phase are called <i>intermediate contexts</i>.</p>

FN\_ctx\_t(3XFN)

Normal resolution of names in context operations always follows XFN links.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `FN_attrset_t(3XFN)`, `FN_composite_name_t(3XFN)`, `FN_ref_t(3XFN)`, `FN_status_t(3XFN)`, `fn_ctx_bind(3XFN)`, `fn_ctx_create_subcontext(3XFN)`, `fn_ctx_destroy_subcontext(3XFN)`, `fn_ctx_get_ref(3XFN)`, `fn_ctx_get_syntax_attrs(3XFN)`, `fn_ctx_handle_destroy(3XFN)`, `fn_ctx_handle_from_initial(3XFN)`, `fn_ctx_handle_from_ref(3XFN)`, `fn_ctx_list_bindings(3XFN)`, `fn_ctx_list_names(3XFN)`, `fn_ctx_lookup(3XFN)`, `fn_ctx_lookup_link(3XFN)`, `fn_ctx_rename(3XFN)`, `fn_ctx_unbind(3XFN)`, `xfn(3XFN)`, `xfn_links(3XFN)`, `xfn_status_codes(3XFN)`, `attributes(5)`

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.



<b>NAME</b>	fn_ctx_unbind – unbind a name from a context				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  int <b>fn_ctx_unbind</b>(FN_ctx_t *ctx, const FN_composite_name_t *name,                   FN_status_t *status);</pre>				
<b>DESCRIPTION</b>	<p>This operation removes the terminal atomic name in <i>name</i> from the the target context — that named by all but the terminal atomic part of <i>name</i>.</p> <p>This operation is successful even if the terminal atomic name was not bound in target context, but fails if any of the intermediate names are not bound. <code>fn_ctx_unbind()</code> is idempotent.</p>				
<b>RETURN VALUE</b>	The operation returns 1 if successful, and 0 otherwise.				
<b>ERRORS</b>	<p><code>fn_ctx_unbind()</code> sets <i>status</i> as described in <code>FN_status_t</code> and <code>xfn_status_codes</code> (3XFN).</p> <p>Certain naming systems may disallow unbinding a name if the name is bound to an existing context in order to avoid orphan contexts that cannot be reached via any name. In such situations, the status code <code>FN_E_OPERATION_NOT_SUPPORTED</code> is returned.</p>				
<b>APPLICATION USAGE</b>	<p>In naming systems that support attributes, and store the attributes along with the names, the unbind operation removes the name and its associated attributes.</p> <p>Normal resolution always follows links. In an <code>fn_ctx_unbind()</code> operation, resolution of <i>name</i> continues to the target context; the terminal atomic name is not resolved. If the terminal atomic name is bound to a link, the link is not followed and the link itself is unbound from the terminal atomic name.</p>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe.</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe.
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe.				
<b>SEE ALSO</b>	<code>FN_composite_name_t(3XFN)</code> , <code>FN_ctx_t(3XFN)</code> , <code>FN_ref_t(3XFN)</code> , <code>FN_status_t(3XFN)</code> , <code>fn_ctx_bind(3XFN)</code> , <code>fn_ctx_lookup(3XFN)</code> , <code>xfn_status_codes(3XFN)</code> , <code>xfn(3XFN)</code> , <code>attributes(5)</code>				

## FN\_identifier\_t(3XFN)

<b>NAME</b>	FN_identifier_t – an XFN identifier								
<b>DESCRIPTION</b>	<p>Identifiers are used to identify reference types and address types in an XFN reference, and to identify attributes and their syntax in the attribute operations.</p> <p>An XFN identifier consists of an unsigned int, which determines the format of identifier, and the actual identifier, which is expressed as a sequence of octets.</p> <p>The representation of this structure is defined by XFN as follows:</p> <pre>typedef struct {     unsigned int format;     size_t length;     void *contents; } FN_identifier_t;</pre> <p>XFN defines a small number of standard forms for identifiers:</p> <table><tr><td>FN_ID_STRING</td><td>The identifier is an ASCII string (ISO 646).</td></tr><tr><td>FN_ID_DCE_UUID</td><td>The identifier is an OSF DCE UUID in string representation. (See the X/Open DCE RPC.)</td></tr><tr><td>FN_ID_ISO_OID_STRING</td><td>The identifier is an ISO OID in ASN.1 dot-separated integer list string format. (See the ISO ASN.1.)</td></tr><tr><td>FN_ID_ISO_OID_BER</td><td>The identifier is an ISO OID in ASN.1 Basic Encoding Rules (BER) format. (See the ISO BER.)</td></tr></table>	FN_ID_STRING	The identifier is an ASCII string (ISO 646).	FN_ID_DCE_UUID	The identifier is an OSF DCE UUID in string representation. (See the X/Open DCE RPC.)	FN_ID_ISO_OID_STRING	The identifier is an ISO OID in ASN.1 dot-separated integer list string format. (See the ISO ASN.1.)	FN_ID_ISO_OID_BER	The identifier is an ISO OID in ASN.1 Basic Encoding Rules (BER) format. (See the ISO BER.)
FN_ID_STRING	The identifier is an ASCII string (ISO 646).								
FN_ID_DCE_UUID	The identifier is an OSF DCE UUID in string representation. (See the X/Open DCE RPC.)								
FN_ID_ISO_OID_STRING	The identifier is an ISO OID in ASN.1 dot-separated integer list string format. (See the ISO ASN.1.)								
FN_ID_ISO_OID_BER	The identifier is an ISO OID in ASN.1 Basic Encoding Rules (BER) format. (See the ISO BER.)								
<b>FILES</b>	#include <xfn/xfn.h>								
<b>SEE ALSO</b>	FN_attribute_t(3XFN), FN_ref_addr_t(3XFN), FN_ref_t(3XFN), xfn(3XFN)								
<b>NOTES</b>	The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.								

<b>NAME</b>	FN_ref_addr_t, fn_ref_addr_create, fn_ref_addr_destroy, fn_ref_addr_copy, fn_ref_addr_assign, fn_ref_addr_type, fn_ref_addr_length, fn_ref_addr_data, fn_ref_addr_description – an address in an XFN reference
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_ref_addr_t *fn_ref_addr_create(constFN_identifier_t *type,     size_t length, const void *data);  void fn_ref_addr_destroy(FN_ref_addr_t *addr);  FN_ref_addr_t *fn_ref_addr_copy(constFN_ref_addr_t *addr);  FN_ref_addr_t *fn_ref_addr_assign(FN_ref_addr_t *dst, const     FN_ref_addr_t *src);  const FN_identifier_t *fn_ref_addr_type(constFN_ref_addr_t     *addr);  size_t fn_ref_addr_length(const FN_ref_addr_t *addr);  const void* fn_ref_addr_data(const FN_ref_addr_t *addr);  FN_string_t *fn_ref_addr_description(constFN_ref_addr_t *addr,     unsigned int detail, unsigned int *more_detail);</pre>
<b>DESCRIPTION</b>	<p>An XFN reference is represented by the type FN_ref_t. An object of this type contains a reference type and a list of addresses. Each address in the list is represented by an object of type FN_ref_addr_t. An address consists of an opaque data buffer and a type field, of type FN_identifier_t.</p> <p>fn_ref_addr_create() creates and returns an address with the given type and data. <i>length</i> indicates the size of the data. fn_ref_addr_destroy() releases the storage associated with the given address. fn_ref_addr_copy() returns a copy of the given address object. fn_ref_addr_assign() makes a copy of the address pointed to by <i>src</i> and assigns it to <i>dst</i>, releasing any old contents of <i>dst</i>. A pointer to the same object as <i>dst</i> is returned.</p> <p>fn_ref_addr_type() returns the type of the given address.  fn_ref_addr_length() returns the size of the address in bytes.  fn_ref_addr_data() returns the contents of the address.</p> <p>fn_ref_addr_description() returns the implementation-defined textual description of the address. It takes as arguments a number, <i>detail</i>, and a pointer to a number, <i>more_detail</i>. <i>detail</i> specifies the level of detail for which the description should be generated; the higher the number, the more detail is to be provided. If <i>more_detail</i> is 0, it is ignored. If <i>more_detail</i> is non-zero, it is set by the description operation to indicate the next level of detail available, beyond that specified by <i>detail</i>. If no higher level of detail is available, <i>more_detail</i> is set to <i>detail</i>.</p>

## FN\_ref\_addr\_t(3XFN)

**USAGE** The address type of an `FN_ref_addr_t` object is intended to identify the mechanism that should be used to reach the object using that address. The client must interpret the contents of the opaque data buffer of the address based on the type of the address, and on the type of the reference that the address is in. However, this interpretation is intended to occur below the application layer. Most applications developers should not have to manipulate the contents of either address or reference objects themselves. These interfaces would generally be used within service libraries.

Multiple addresses in a single reference are intended to identify multiple communication endpoints for the same conceptual object. Multiple addresses may arise for various reasons, such as the object offering interfaces over more than one communication mechanism.

Manipulation of addresses using the operations described in this manual page does not affect their representation in the underlying naming system. Changes to addresses in the underlying naming system can only be effected through the use of the interfaces described in `FN_ctx_t(3XFN)`.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `FN_ctx_t(3XFN)`, `FN_identifier_t(3XFN)`, `FN_ref_t(3XFN)`, `FN_string_t(3XFN)`, `xfn(3XFN)`, `attributes(5)`

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

<b>NAME</b>	FN_ref_t, fn_ref_create, fn_ref_destroy, fn_ref_copy, fn_ref_assign, fn_ref_type, fn_ref_addrcount, fn_ref_first, fn_ref_next, fn_ref_prepend_addr, fn_ref_append_addr, fn_ref_insert_addr, fn_ref_delete_addr, fn_ref_delete_all, fn_ref_create_link, fn_ref_is_link, fn_ref_link_name, fn_ref_description – an XFN reference
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_ref_t *fn_ref_create(const FN_identifier_t *ref_type); void fn_ref_destroy(FN_ref_t *ref); FN_ref_t *fn_ref_copy(const FN_ref_t *ref); FN_ref_t *fn_ref_assign(FN_ref_t *dst, const FN_ref_t *src); const FN_identifier_t *fn_ref_type(const FN_ref_t *ref); unsigned int fn_ref_addrcount(const FN_ref_t *ref); const FN_ref_addr_t *fn_ref_first(const FN_ref_t *ref, void **iter_pos); const FN_ref_addr_t *fn_ref_next(const FN_ref_t *ref, void **iter_pos); int fn_ref_prepend_addr(FN_ref_t *ref, const FN_ref_addr_t *addr); int fn_ref_append_addr(FN_ref_t *ref, const FN_ref_addr_t *addr); int fn_ref_insert_addr(FN_ref_t *ref, void **iter_pos, const FN_ref_addr_t *addr); int fn_ref_delete_addr(FN_ref_t *ref, void **iter_pos); int fn_ref_delete_all(FN_ref_t *ref); FN_ref_t *fn_ref_create_link(const FN_composite_name_t *link_name); int fn_ref_is_link(const FN_ref_t *ref); FN_composite_name_t *fn_ref_link_name(const FN_ref_t *link_ref); FN_string_t *fn_ref_description(const FN_ref_t *ref, unsigned int detail, unsigned int *more_detail);</pre>
<b>DESCRIPTION</b>	<p>An XFN reference is represented by the type FN_ref_t. An object of this type contains a reference type and a list of addresses. The ordering in this list at the time of binding might not be preserved when the reference is returned upon lookup.</p> <p>The reference type is represented by an object of type FN_identifier_t. The reference type is intended to identify the class of object referenced. XFN does not dictate the precise use of this.</p> <p>Each address is represented by an object of type FN_ref_addr_t.</p>

## FN\_ref\_t(3XFN)

`fn_ref_create()` creates a reference with no address, using `ref_type` as its reference type. Addresses can be added later to the reference using the functions described below. `fn_ref_destroy()` releases the storage associated with `ref`. `fn_ref_copy()` creates a copy of `ref` and returns it. `fn_ref_assign()` creates a copy of `src` and assigns it to `dst`, releasing any old contents of `dst`. A pointer to the same object as `dst` is returned.

`fn_ref_addrcount()` returns the number of addresses in the reference `ref`.

`fn_ref_first()` returns the first address in `ref` and sets `iter_pos` to be after the address. It returns 0 if there is no address in the list. `fn_ref_next()` returns the address following `iter_pos` in `ref` and sets `iter_pos` to be after the address. If the iteration marker `iter_pos` is at the end of the sequence, `fn_ref_next()` returns 0.

`fn_ref_prepend_addr()` adds `addr` to the front of the list of addresses in `ref`.

`fn_ref_append_addr()` adds `addr` to the end of the list of addresses in `ref`.

`fn_ref_insert_addr()` adds `addr` to `ref` before `iter_pos` and sets `iter_pos` to be immediately after the new reference added. `fn_ref_delete_addr()` deletes the address located before `iter_pos` in the list of addresses in `ref` and sets `iter_pos` back one address. `fn_ref_delete_all()` deletes all addresses in `ref`.

`fn_ref_create_link()` creates a reference using the given composite name `link_name` as an address. `fn_ref_is_link()` tests if `ref` is a link. It returns 1 if it is; 0 if it is not. `fn_ref_link_name()` returns the composite name stored in a link reference. It returns 0 if `link_ref` is not a link.

`fn_ref_description()` returns a string description of the given reference. It takes as argument an integer, `detail`, and a pointer to an integer, `more_detail`. `detail` specifies the level of detail for which the description should be generated; the higher the number, the more detail is to be provided. If `more_detail` is 0, it is ignored. If `more_detail` is non-zero, it is set by the description operation to indicate the next level of detail available, beyond that specified by `detail`. If no higher level of detail is available, `more_detail` is set to `detail`.

### RETURN VALUES

The following operations return 1 if the operation succeeds, 0 if the operation fails:

```
fn_ref_prepend_addr()
fn_ref_append_addr()
fn_ref_insert_addr()
fn_ref_delete_addr()
fn_ref_delete_all()
```

### USAGE

The reference type is intended to identify the class of object referenced. XFN does not dictate the precise use of this.

Multiple addresses in a single reference are intended to identify multiple communication endpoints for the same conceptual object. Multiple addresses may arise for various reasons, such as the object offering interfaces over more than one communication mechanism.

The client must interpret the contents of a reference based on the type of the addresses and the type of the reference. However, this interpretation is intended to occur below the application layer. Most applications developers should not have to manipulate the contents of either address or reference objects themselves. These interfaces would generally be used within service libraries.

Manipulation of references using the operations described in this manual page does not affect their representation in the underlying naming system. Changes to references in the underlying naming system can only be effected through the use of the interfaces described in FN\_ctx\_t(3XFN).

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_composite\_name\_t(3XFN), FN\_ctx\_t(3XFN), FN\_identifier\_t(3XFN), FN\_ref\_addr\_t(3XFN), FN\_string\_t(3XFN), fn\_ctx\_lookup(3XFN), fn\_ctx\_lookup\_link(3XFN), xfn(3XFN), xfn\_links(3XFN), `attributes(5)`

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## FN\_search\_control\_t(3XFN)

<b>NAME</b>	FN_search_control_t, fn_search_control_create, fn_search_control_destroy, fn_search_control_copy, fn_search_control_assign, fn_search_control_scope, fn_search_control_follow_links, fn_search_control_max_names, fn_search_control_return_ref, fn_search_control_return_attr_ids – options for attribute search						
<b>SYNOPSIS</b>	<pre>#include &lt;xfn/xfn.h&gt;  FN_search_control_t *fn_search_control_create(unsigned int scope,         unsigned int follow_links, unsigned int max_names, unsigned int         return_ref, const FN_attrset_t *return_attr_ids, unsigned int         *status);  void fn_search_control_destroy(FN_search_control_t *scontrol);  FN_search_control_t *fn_search_control_copy(const         FN_search_control_t *scontrol);  FN_search_control_t         *fn_search_control_assign(FN_search_control_t *dst, const         FN_search_control_t *src);  unsigned int fn_search_control_scope(const FN_search_control_t         *scontrol);  unsigned int fn_search_control_follow_links(const         FN_search_control_t *scontrol);  unsigned int fn_search_control_max_names(const         FN_search_control_t *scontrol);  unsigned int fn_search_control_return_ref(const         FN_search_control_t *scontrol);  const FN_attrset_t *fn_search_control_return_attr_ids(const         FN_search_control_t *scontrol);</pre>						
<b>DESCRIPTION</b>	<p>The FN_search_control_t object is used to specify options for the attribute search operation fn_attr_ext_search(3XFN).</p> <p>fn_search_control_create() creates an FN_search_control_t object using information in <i>scope</i>, <i>follow_links</i>, <i>max_names</i>, <i>return_ref</i>, and <i>return_attr_ids</i> to set the search options. If the operation succeeds, fn_search_control_create() returns a pointer to an FN_search_control_t object; otherwise, it returns a NULL pointer.</p> <p>The scope of the search, <i>scope</i>, is either the named object, the named context, the named context and its subcontexts, or the named context and a context implementation defined set of subcontexts. The values for <i>scope</i> are:</p> <table><tr><td>FN_SEARCH_NAMED_OBJECT</td><td>Search just the given named object.</td></tr><tr><td>FN_SEARCH_ONE_CONTEXT</td><td>Search just the given context.</td></tr><tr><td>FN_SEARCH_SUBTREE</td><td>Search given context and all its subcontexts.</td></tr></table>	FN_SEARCH_NAMED_OBJECT	Search just the given named object.	FN_SEARCH_ONE_CONTEXT	Search just the given context.	FN_SEARCH_SUBTREE	Search given context and all its subcontexts.
FN_SEARCH_NAMED_OBJECT	Search just the given named object.						
FN_SEARCH_ONE_CONTEXT	Search just the given context.						
FN_SEARCH_SUBTREE	Search given context and all its subcontexts.						



## FN\_search\_control\_t(3XFN)

`FN_SEARCH_CONSTRAINED_SUBTREE` Search given context and its subcontexts as constrained by the context-specific policy in place at the named context.

*follow\_links* further defines the scope and nature of the search. If *follow\_links* is nonzero, the search follows XFN links. If *follow\_links* is 0, XFN links are not followed. See `fn_attr_ext_search(3XFN)` for more detail about how XFN links are treated.

*max\_names* specifies the maximum number of names to return in an `FN_ext_searchlist_t(3XFN)` enumeration (see `fn_attr_ext_search(3XFN)`). The names of all objects whose attributes satisfy the filter are returned when *max\_names* is 0.

If *return\_ref* is non-zero, the reference bound to the named object is returned with the object's name by `fn_ext_searchlist_next(3XFN)` (see `fn_attr_ext_search(3XFN)`). If *return\_ref* is 0, the reference is not returned.

Attribute identifiers and values associated with named objects that satisfy the filter may be returned by `fn_ext_searchlist_next(3XFN)`. The attributes returned are those listed in *return\_attr\_ids*. If the value of *return\_attr\_ids* is 0, all attributes are returned. If *return\_attr\_ids* is an empty `FN_attrset_t` object (see `FN_attrset_t(3XFN)`), no attributes are returned. Any attribute values in *return\_attr\_ids* are ignored; only the attribute identifiers are relevant for this operation.

`fn_attr_ext_search(3XFN)` interprets a value of 0 for the search control argument as a default search control which has the following option settings:

<i>scope</i>	<code>FN_SEARCH_ONE_CONTEXT</code>
<i>follow_links</i>	0 (do not follow links)
<i>max_names</i>	0 (return all named objects that match filter)
<i>return_ref</i>	0 (do not return the reference of the named object)
<i>return_attr_ids</i>	an empty <code>FN_attrset_t</code> object (do not return any attributes of the named object)

`fn_search_control_destroy()` releases the storage associated with *scontrol*.

`fn_search_control_copy()` returns a copy of the search control *scontrol*.

`fn_search_control_assign()` makes a copy of the search control *src* and assigns it to *dst*, releasing the old contents of *dst*. A pointer to the same object as *dst* is returned.

`fn_search_control_scope()` returns the scope for the search.

`fn_search_control_follow_links()` returns non-zero if links are followed; 0 if not.

`fn_search_control_max_names()` returns the maximum number of names.

## FN\_search\_control\_t(3XFN)

fn\_search\_control\_return\_ref() returns nonzero if the reference is returned; 0 if not.

fn\_search\_control\_return\_attr\_ids() returns a pointer to the list of attributes; a NULL pointer indicates that all attributes and values are returned.

**ERRORS** fn\_search\_control\_create() returns a NULL pointer if the operation fails and sets status as follows:

FN\_E\_SEARCH\_INVALID\_OPTION      A supplied search option was invalid or inconsistent.

Other status codes are possible (see xfn\_status\_codes(3XFN)).

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_attrset\_t(3XFN), fn\_attr\_ext\_search(3XFN), xfn\_status\_codes(3XFN), attributes(5)

<b>NAME</b>	FN_search_filter_t, fn_search_filter_create, fn_search_filter_destroy, fn_search_filter_copy, fn_search_filter_assign, fn_search_filter_expression, fn_search_filter_arguments – filter expression for attribute search
<b>SYNOPSIS</b>	<pre>#include &lt;xfn/xfn.h&gt;  FN_search_filter_t *fn_search_filter_create(unsigned int *status,       const unsigned char *estr, .);  void fn_search_filter_destroy(FN_search_filter_t *sfilter);  FN_search_filter_t *fn_search_filter_copy(const       FN_search_filter_t *sfilter);  FN_search_filter_t *fn_search_filter_assign(FN_search_filter_t       *dst, const FN_search_filter_t *src);  const char *fn_search_filter_expression(const FN_search_filter_t       *sfilter);  const void **fn_search_filter_arguments(const FN_search_filter_t       *sfilter, size_t *number_of_arguments);</pre>
<b>DESCRIPTION</b>	<p>The <code>FN_search_filter_t</code> type is an expression that is evaluated against the attributes of named objects bound in the scope of the search operation <code>fn_attr_ext_search(3XFN)</code>. The filter evaluates to <code>TRUE</code> or <code>FALSE</code>. If the filter is empty, it evaluates to <code>TRUE</code>. Names of objects whose attribute values satisfy the filter expression are returned by the search operation.</p> <p>If the identifier in any subexpression of the filter does not exist as an attribute of an object, then the innermost logical expression containing that identifier is <code>FALSE</code>. A subexpression that is only an attribute tests for the presence of the attribute; the subexpression evaluates to <code>TRUE</code> if the attribute has been defined for the object and <code>FALSE</code> otherwise.</p> <p><code>fn_search_filter_create()</code> creates a search filter from the expression string <i>estr</i> and the remaining arguments.</p> <p><code>fn_search_filter_destroy()</code> releases the storage associated with the search filter <i>sfilter</i>.</p> <p><code>fn_search_filter_copy()</code> returns a copy of the search filter <i>sfilter</i>.</p> <p><code>fn_search_filter_assign()</code> makes a copy of the search filter <i>src</i> and assigns it to <i>dst</i>, releasing the old contents of <i>dst</i>. A pointer to the same object as <i>dst</i> is returned.</p> <p><code>fn_search_filter_expression()</code> returns the filter expression of <i>sfilter</i>.</p> <p><code>fn_search_filter_arguments()</code> returns an array of pointers to arguments supplied to the filter constructor. <i>number_of_arguments</i> is set to the size of this array. The types of the arguments are determined by the substitution tokens in the expression in <i>sfilter</i>.</p>

## FN\_search\_filter\_t(3XFN)

### BNF of Filter Expression

```

<FilterExpr> : : = [ <Expr> ]
<Expr> : : = <Expr> "or" <Expr>
          <Expr> "and" <Expr>
          | "not" <Expr>
          | "(" <Expr> ")"
          | <Attribute> [ <Rel_Op> <Value> ]
          | <Ext>
<Rel_Op> : : = "==" | "!=" | "<" | "< =" | ">" | "> =" | "≈ ="
<Attribute> : : = "%a"
<Value> : : = <Integer>
            | "%v"
            | <Wildcarded_string>
<Wildcarded_string> : : = "*"
                    | <String>
                    | {<String> "*" }+ [<String>]
                    | {"*" <String>}+ ["*"]
<String> : : = "`" { <Char> } * "`"
            | "%s"
<Char> : : = <PCS> // See BNF in Section 4.1.2 for PCSdefinition
            | Characters in the repertoire of a string representation
<Identifier> : : = "%i"
<Ext> : : = <Ext_Op> "(" [Arg_List] ")"
<Ext_Op> : : = <String> | <Identifier>
<Arg_List> : : = <Arg> | <Arg> "," <Arg_List>
<Arg> : : = <Value> | <Attribute> | <Identifier>

```

### Specification of Filter Expression

The arguments to `fn_search_filter_create()` are a return status, an expression string, and a list of arguments. The string contains the filter expression with substitution tokens for the attributes, attribute values, strings, and identifiers that are part of the expression. The remaining list of arguments contains the attributes and values in the order of appearance of their corresponding substitution tokens in the expression. The arguments are of types `FN_attribute_t*`, `FN_attrvalue_t*`, `FN_string_t*`, or `FN_identifier_t*`. Any attribute values in an `FN_attribute_t*` type of argument are ignored; only the attribute identifier and attribute syntax are relevant. The argument type expected by each substitution token are listed in the following table.

Token	Argument Type
%a	FN_attribute_t*
%v	FN_attrvalue_t*
%s	FN_string_t*
%i	FN_identifier_t*

### Precedence

The following precedence relations hold in the absence of parentheses, in the order of lowest to highest:

or  
and

**Relational Operators**

not  
relational operators

These boolean and relational operators are left associative.

Comparisons and ordering are specific to the syntax and/or rules of the supplied attribute.

Locale (code set, language, or territory) mismatches that occur during string comparisons and ordering operations are resolved in an implementation-dependent way. Relational operations that have ordering semantics may be used for strings of code sets in which ordering is meaningful, but is not of general use in internationalized environments.

An attribute that occurs in the absence of any relational operator tests for the presence of the attribute.

Operator	Meaning
==	The sub-expression is TRUE if at least one value of the specified attribute is equal to the supplied value.
! =	The sub-expression is TRUE if no values of the specified attribute equal the supplied value.
> =	The sub-expression is TRUE if at least one value of the attribute is greater than or equal to the supplied value.
>	The sub-expression is TRUE if at least one value of the attribute is greater than the supplied value.
< =	The sub-expression is TRUE if at least one value of the attribute is less than or equal to the supplied value.
<	The sub-expression is TRUE if at least one value of the attribute is less than the supplied value.
≈ =	The sub-expression is TRUE if at least one value of the specified attribute matches the supplied value according to some context-specific approximate matching criterion. This criterion must subsume strict equality.

**Wildcarded Strings**

A wildcarded string consists of a sequence of alternating wildcard specifiers and strings. The sequence can start with either a wildcard specifier or a string, and end with either a wildcard specifier or a string.

The wildcard specifier is denoted by the asterisk character ('\*') and means zero or more occurrences of any character.

Wildcarded strings can be used to specify substring matches. The following are examples of wildcarded strings and what they mean:

## FN\_search\_filter\_t(3XFN)

Wildcarded String	Meaning
*	Any string
*'ing'	Any string ending with ing
Any string starting with jo, and containing the substring ph, and which contains the substring ne in the portion of the string following ph, and which ends with er	
T}	
%s*	Any string starting with the supplied string
Any string starting with bix and ending with the supplied string	
T}	

String matches involving strings of different locales (code set, language, or territory) are resolved in an implementation-dependent way.

### Extended Operations

In addition to the relational operators, extended operators can be specified. All extended operators return either TRUE or FALSE. A filter expression can contain both relational and extended operations.

Extended operators are specified using an identifier (see `FN_identifier_t(3XFN)`) or a string. If the operator is specified using a string, the string is used to construct an identifier of format `FN_ID_STRING`. Identifiers of extended operators and signatures of the corresponding extended operations, as well as their suggested semantics, are registered with X/Open Company Ltd.

The following three extended operations are currently defined:

'name' (<Wildcarded String>)	The identifier for this operation is 'name' ( <code>FN_ID_STRING</code> ). The argument to this operation is a wildcard string. The operation returns TRUE if the name of the object matches the supplied wildcard string.
'reftype' (%i)	The identifier for this operation is 'reftype' ( <code>FN_ID_STRING</code> ). The argument to this operation is an identifier.

## FN\_search\_filter\_t(3XFN)

The operation returns TRUE if the reference type of the object is equal to the supplied identifier.

'addrtype' (%i)

The identifier for this operation is 'addrtype' (LM FN\_ID\_STRING). The argument to the operation is an identifier. The operation returns TRUE if any of the address types in the reference of the object is equal to the supplied identifier.

Support and exact semantics of extended operations are context-specific. If a context does not support an extended operation, or if the filter expression supplies the extended operation with either an incorrect number or type of arguments, the error FN\_E\_SEARCH\_INVALID\_OP is returned. (Note: FN\_E\_OPERATION\_NOT\_SUPPORTED is returned when fn\_attr\_ext\_search(3XFN) is not supported.)

The following are examples of filter expressions that contain extended operations:

Expression	Meaning
Evaluates to TRUE if the name of the object starts with bill. T}	
%i(%a, %v)	Evaluates to result of applying the specified operation to the supplied arguments.
(%a == %v) and 'name' ('joe'*)	Evaluates to TRUE if the specified attribute has the given value and if the name of the object starts with joe.

### RETURN VALUES

fn\_search\_filter\_create() returns a pointer to an FN\_search\_filter\_t object if the operation succeeds; otherwise it returns a NULL pointer.

### ERRORS

fn\_search\_filter\_create() returns a NULL pointer if the operation fails and sets *status* in the following way:

FN\_E\_SEARCH\_INVALID\_FILTER

The filter expression had a syntax error or some other problem.

FN\_E\_SEARCH\_INVALID\_OP

An operator in the filter expression is not supported or, if the operator is an extended

## FN\_search\_filter\_t(3XFN)

operator, the number of types of arguments supplied does not match the signature of the operation.

FN\_E\_INVALID\_ATTR\_IDENTIFIER The left hand side of an operator expression was not an attribute.

FN\_E\_INVALID\_ATTR\_VALUE The right hand side of an operator expression was not an integer, attribute value, or (wildcarded) string.

Other status codes are possible as described in the reference manual pages for FN\_status\_t(3XFN) and xfn\_status\_codes(3XFN).

### EXAMPLES

#### EXAMPLE 1 Creating Different Filters

The following examples illustrate how to create three different filters.

The first example shows how to construct a filter involving substitution tokens and literals in the same filter expression. This example creates a filter for named objects whose `color` attribute contains a string value of `red`, `blue`, or `white`. The first two values are specified using substitution tokens; the last value, `white`, is specified as a literal in the expression.

```
unsigned int status;
extern FN_attribute_t *attr_color;
FN_string_t *red = fn_string_from_str((unsigned char *)"red");
FN_string_t *blue = fn_string_from_str((unsigned char *)"blue");
FN_search_filter_t *sfilter;
sfilter = fn_search_filter_create(
    &status,
    "(%a == %s) or (%a == %s) or (%a == 'white')",
    attr_color, red, attr_color, blue,
    attr_color);
```

The second example illustrates how to construct a filter involving a wildcarded string. This example creates a filter for searching for named objects whose `last_name` attribute has a value that begins with the character `m`.

```
unsigned int status;
extern FN_attribute_t *attr_last_name;
FN_search_filter_t *sfilter;
sfilter = fn_search_filter_create(
    &status, "%a == 'm'", attr_last_name);
```

The third example illustrates how to construct a filter involving extended operations. This example creates a filter for finding all named objects whose name ends with `ton`.

```
unsigned int status;
FN_search_filter_t *sfilter;
sfilter = fn_search_filter_create(&status, "'name' (*'ton')");
```

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:



FN\_search\_filter\_t(3XFN)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_attribute\_t(3XFN), FN\_attrvalue\_t(3XFN), FN\_identifier\_t(3XFN), FN\_status\_t(3XFN), FN\_string\_t(3XFN), fn\_attr\_ext\_search(3XFN), xfn\_status\_codes(3XFN), attributes(5)

## FN\_status\_t(3XFN)

<b>NAME</b>	FN_status_t, fn_status_create, fn_status_destroy, fn_status_copy, fn_status_assign, fn_status_code, fn_status_remaining_name, fn_status_resolved_name, fn_status_resolved_ref, fn_status_diagnostic_message, fn_status_link_code, fn_status_link_remaining_name, fn_status_link_resolved_name, fn_status_link_resolved_ref, fn_status_link_diagnostic_message, fn_status_is_success, fn_status_set_success, fn_status_set, fn_status_set_code, fn_status_set_remaining_name, fn_status_set_resolved_name, fn_status_set_resolved_ref, fn_status_set_diagnostic_message, fn_status_set_link_code, fn_status_set_link_remaining_name, fn_status_set_link_resolved_name, fn_status_set_link_resolved_ref, fn_status_set_link_diagnostic_message, fn_status_append_resolved_name, fn_status_append_remaining_name, fn_status_advance_by_name, fn_status_description – an XFN status object
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxfn [ library ... ] #include &lt;xfn/xfn.h&gt;  FN_status_t *fn_status_create(void);  void fn_status_destroy(FN_status_t *stat);  FN_status_t *fn_status_copy(const FN_status_t *stat);  FN_status_t *fn_status_assign(FN_status_t *dst, const FN_status_t     *src);  unsigned int fn_status_code(const FN_status_t *stat);  const FN_composite_name_t     *fn_status_remaining_name(const FN_status_t *stat);  const FN_composite_name_t     *fn_status_resolved_name(const FN_status_t *stat);  const FN_ref_t *fn_status_resolved_ref(const FN_status_t *stat);  const FN_string_t *fn_status_diagnostic_message(const FN_status_t     *stat);  unsigned int fn_status_link_code(const FN_status_t *stat);  const FN_composite_name_t     *fn_status_link_remaining_name(const FN_status_t *stat);  const FN_composite_name_t     *fn_status_link_resolved_name(const FN_status_t *stat);  const FN_ref_t *fn_status_link_resolved_ref(const FN_status_t     *stat);  const FN_string_t     *fn_status_link_diagnostic_message(const FN_status_t *stat);  int fn_status_is_success(const FN_status_t *stat);  int fn_status_set_success(FN_status_t *stat);</pre>

FN\_status\_t(3XFN)

```
int fn_status_set(FN_status_t *stat, unsigned int code, const
    FN_ref_t *resolved_ref, const FN_composite_name_t *resolved_name,
    const FN_composite_name_t *remaining_name);

int fn_status_set_code(FN_status_t *stat, unsigned int code);

int fn_status_set_remaining_name(FN_status_t *stat, const
    FN_composite_name_t *name);

int fn_status_set_resolved_name(FN_status_t *stat, const
    FN_composite_name_t *name);

int fn_status_set_resolved_ref(FN_status_t *stat, const FN_ref_t
    *ref);

int fn_status_set_diagnostic_message(FN_status_t *stat, const
    FN_string_t *msg);

int fn_status_set_link_code(FN_status_t *stat, unsigned int code);

int fn_status_set_link_remaining_name(FN_status_t *stat, const
    FN_composite_name_t *name);

int fn_status_set_link_resolved_name(FN_status_t *stat, const
    FN_composite_name_t *name);

int fn_status_set_link_resolved_ref(FN_status_t *stat, const
    FN_ref_t *ref);

int fn_status_set_link_diagnostic_message(FN_status_t *stat, const
    FN_string_t *msg);

int fn_status_append_resolved_name(FN_status_t *stat, const
    FN_composite_name_t *name);

int fn_status_append_remaining_name(FN_status_t *stat, const
    FN_composite_name_t *name);

int fn_status_advance_by_name(FN_status_t *stat, const
    FN_composite_name_t *prefix, const FN_ref_t *resolved_ref);

FN_string_t *fn_status_description(const FN_status_t *stat,
    unsigned int detail, unsigned int *more_detail);
```

## DESCRIPTION

The result status of operations in the context interface and the attribute interface is encapsulated in an `FN_status_t` object. This object contains information about how the operation completed: whether an error occurred in performing the operation, the nature of the error, and information that helps locate where the error occurred. In the case that the error occurred while resolving an XFN link, the status object contains additional information about that error.

The context status object consists of several items of information:

primary status code	An unsigned int code describing the disposition of the operation.
---------------------	---

## FN\_status\_t(3XFN)

resolved name	In the case of a failure during the resolution phase of the operation, this is the leading portion of the name that was resolved successfully. Resolution may have been successful beyond this point, but the error might not be pinpointed further.
resolved reference	The reference to which resolution was successful (in other words, the reference to which the resolved name is bound).
remaining name	The remaining unresolved portion of the name.
diagnostic message	This contains any diagnostic message returned by the context implementation. This message provides the context implementation a way of notifying the end-user or administrator of any implementation-specific information related to the returned error status. The diagnostic message could then be used by the end-user or administrator to take appropriate out-of-band action to rectify the problem.
link status code	In the case that an error occurred while resolving an XFN link, the primary status code has the value <code>FN_E_LINK_ERROR</code> and the link status code describes the error that occurred while resolving the XFN link.
resolved link name	In the case of a link error, this contains the resolved portion of the name in the XFN link.
resolved link reference	In the case of a link error, this contains the reference to which the resolved link name is bound.
remaining link name	In the case of a link error, this contains the remaining unresolved portion of the name in the XFN link.
link diagnostic message	In the case of a link error, this contains any diagnostic message related to the resolution of the link.

Both the primary status code and the link status code are values of type `unsigned int` that are drawn from the same set of meaningful values. XFN reserves the values 0 through 127 for standard meanings. The values and interpretations for the codes are determined by XFN. See `xfn_status_codes(3XFN)`.

`fn_status_create()` creates a status object with status `FN_SUCCESS`.

`fn_status_destroy()` releases the storage associated with *stat*.

`fn_status_copy()` returns a copy of the status object *stat*. `fn_status_assign()` makes a copy of the status object *src* and assigns it to *dst*, releasing any old contents of *dst*. A pointer to the same object as *dst* is returned.

`fn_status_code()` returns the status code. `fn_status_remaining_name()` returns the remaining part of name to be resolved. `fn_status_resolved_name()` returns the part of the composite name that has been resolved. `fn_status_resolved_ref()` returns the reference to which resolution was successful. `fn_status_diagnostic_message` returns any diagnostic message set by the context implementation.

`fn_status_link_code()` returns the link status code. `fn_status_link_remaining_name()` returns the remaining part of the link name that has not been resolved. `fn_status_link_resolved_name()` returns the part of the link name that has been resolved. `fn_status_link_resolved_ref()` returns the reference to which resolution of the link was successful. `fn_status_link_diagnostic_message()` returns any diagnostic message set by the context implementation during resolution of the link.

`fn_status_is_success()` returns 1 if the status indicates success, 0 otherwise.

`fn_status_set_success()` sets the status code to `FN_SUCCESS` and clears all other parts of *stat*. `fn_status_set()` sets the non-link contents of the status object *stat*. `fn_status_set_code()` sets the primary status code field of the status object *stat*. `fn_status_set_remaining_name()` sets the remaining name part of the status object *stat* to *name*. `fn_status_set_resolved_name()` sets the resolved name part of the status object *stat* to *name*. `fn_status_set_resolved_ref()` sets the resolved reference part of the status object *stat* to *ref*. `fn_status_set_diagnostic_message()` sets the diagnostic message part of the status object to *msg*.

`fn_status_set_link_code()` sets the link status code field of the status object *stat* to indicate why resolution of the link failed.

`fn_status_set_link_remaining_name()` sets the remaining link name part of the status object *stat* to *name*. `fn_status_set_link_resolved_name()` sets the resolved link name part of the status object *stat* to *name*.

`fn_status_set_link_resolved_ref()` sets the resolved link reference part of the status object *stat* to *ref*. `fn_status_set_link_diagnostic_message()` sets the link diagnostic message part of the status object to *msg*.

`fn_status_append_resolved_name()` appends as additional components *name* to the resolved name part of the status object *stat*.

`fn_status_append_remaining_name()` appends as additional components *name* to the remaining name part of the status object *stat*.

`fn_status_advance_by_name()` removes *prefix* from the remaining name, and appends it to the resolved name. The resolved reference part is set to *resolved\_ref*. This operation returns 1 on success, 0 if the *prefix* is not a prefix of the remaining name.

**RETURN VALUES**

The `fn_status_set_*( )` operations return 1 if the operation succeeds, 0 if the operation fails.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

## FN\_status\_t(3XFN)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** FN\_composite\_name\_t(3XFN), FN\_ref\_t(3XFN), FN\_string\_t(3XFN), xfn(3XFN), xfn\_status\_codes(3XFN), attributes(5)

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

**NAME** FN\_string\_t, fn\_string\_create, fn\_string\_destroy, fn\_string\_from\_str, fn\_string\_from\_str\_n, fn\_string\_str, fn\_string\_from\_contents, fn\_string\_code\_set, fn\_string\_charcount, fn\_string\_bytcount, fn\_string\_contents, fn\_string\_copy, fn\_string\_assign, fn\_string\_from\_strings, fn\_string\_from\_substring, fn\_string\_is\_empty, fn\_string\_compare, fn\_string\_compare\_substring, fn\_string\_next\_substring, fn\_string\_prev\_substring – a character string

**SYNOPSIS**

```
cc [ flag ... ] file ... -lxfn [ library ... ]
#include <xfn/xfn.h>

FN_string_t *fn_string_create(void);

void fn_string_destroy(FN_string_t *str);

FN_string_t *fn_string_from_str(const unsigned char *cstr);

FN_string_t *fn_string_from_str_n(const unsigned char *cstr,
    size_t n);

const unsigned char *fn_string_str(const FN_string_t *str,
    unsigned int *status);

FN_string_t *fn_string_from_contents(unsigned long code_set, const
    void *locale_info, size_t locale_info_len, size_t charcount, size_t
    bytcount, const void *contents, unsigned int *status);

unsigned long fn_string_code_set(const FN_string_t *str, const
    void **locale_info, size_t *locale_info_len);

size_t fn_string_charcount(const FN_string_t *str);

size_t fn_string_bytcount(const FN_string_t *str);

const void *fn_string_contents(const FN_string_t *str);

FN_string_t *fn_string_copy(const FN_string_t *str);

FN_string_t *fn_string_assign(FN_string_t *dst, const FN_string_t
    *src);

FN_string_t *fn_string_from_strings(unsigned int *status, const
    FN_string_t *s1, const FN_string_t *s2, ...);

FN_string_t *fn_string_from_substring(const FN_string_t *str, int
    first, int last);

int fn_string_is_empty(const FN_string_t *str);

int fn_string_compare(const FN_string_t *str1, const FN_string_t
    *str2, unsigned int string_case, unsigned int *status);

int fn_string_compare_substring(const FN_string_t *str1, int first,
    int last, const FN_string_t *str2, unsigned int string_case,
    unsigned int *status);
```

## FN\_string\_t(3XFN)

```
int fn_string_next_substring(const FN_string_t *str, const
    FN_string_t *sub, int index, unsigned int string_case, unsigned
    int *status);

int fn_string_prev_substring(const FN_string_t *str, const
    FN_string_t *sub, int index, unsigned int string_case, unsigned
    int *status);
```

### DESCRIPTION

The `FN_string_t` type is used to represent character strings in the XFN interface. It provides insulation from specific string representations.

The `FN_string_t` supports multiple code sets. It provides creation functions for character strings of the code set of the current locale setting and a generic creation function for arbitrary code sets. The degree of support for the functions that manipulate `FN_string_t` for arbitrary code sets is implementation-dependent. An XFN implementation is required to support the ISO 646 code set; all other code sets are optional.

`fn_string_destroy()` releases the storage associated with the given string.

`fn_string_create()` creates an empty string.

`fn_string_from_str()` creates an `FN_string_t` object from the given null terminated string based on the code set of the current locale setting. The number of characters in the string is determined by the code set of the current locale setting.

`fn_string_from_str_n()` is like `fn_string_from_str()` except only *n* characters from the given string are used. `fn_string_str()` returns the contents of the given string *str* in the form of a null terminated string in the code set and current locale setting.

`fn_string_from_contents()` creates an `FN_string_t` object using the specified code set *code\_set*, locale information *locale\_info*, and data in the given buffer *contents*. *bytecount* specifies the number of bytes in *contents* and *charcount* specifies the number of characters represented by *contents*.

`fn_string_code_set()` returns the code set associated with the given string object and, if present, the locale information in *locale\_info*. `fn_string_charcount()` returns the number of characters in the given string object.

`fn_string_bytecount()` returns the number of bytes used to represent the given string object. `fn_string_contents()` returns a pointer to the contents of the given string object.

`fn_string_copy()` returns a copy of the given string object.

`fn_string_assign()` makes a copy of the string object *src* and assigns it to *dst*, releasing any old contents of *dst*. A pointer to the same object as *dst* is returned.

`fn_string_from_strings()` is a function that takes a variable number of arguments (minimum of 2), the last of which must be `NULL (0)`; it returns a new string object composed of the left to right concatenation of the given strings, in the given order. The support for strings with different code sets and/or locales as arguments to a single invocation of `fn_string_from_strings()` is implementation-dependent.



`fn_string_from_substring()` returns a new string object consisting of the characters located between *first* and *last* inclusive from *str*. Indexing begins with 0. If *last* is `FN_STRING_INDEX_LAST` or exceeds the length of the string, the index of the last character of the string is used.

`fn_string_is_empty()` returns whether *str* is an empty string.

Comparison of two strings must take into account code set and locale information. If strings are in the same code set and same locale, case sensitivity is applied according to the case sensitivity rules applicable for the code set and locale; case sensitivity may not necessarily be relevant for all string encodings. If *string\_case* is non-zero, case is significant and equality for strings of the same code set is defined as equality between byte-wise encoded values of the strings. If *string\_case* is zero, case is ignored and equality for strings of the same code set is defined using the definition of case-insensitive equality for the specific code set. Support for comparison between strings of different code sets, or lack thereof, is implementation-dependent.

`fn_string_compare()` compares strings *str1* and *str2* and returns 0 if they are equal, non-zero if they are not equal. If two strings are not equal, `fn_string_compare()` returns a positive value if the difference of *str2* precedes that of *str1* in terms of byte-wise encoded value (with case-sensitivity taken into account when *string\_case* is non-zero), and a negative value if the difference of *str1* precedes that of *str2*, in terms of byte-wise encoded value (with case-sensitivity taken into account when *string\_case* is non-zero). Such information (positive versus negative return value) may be used by applications that use strings of code sets in which ordering is meaningful; this information is not of general use in internationalized environments. `fn_string_compare_substring()` is similar to `fn_string_compare()` except that `fn_string_compare_substring()` compares characters between *first* and *last* inclusive of *str2* with *str1*. Comparison of strings with incompatible code sets returns a negative or positive value (never 0) depending on the implementation.

`fn_string_next_substring()` returns the index of the next occurrence of *sub* at or after *index* in the string *str*. `FN_STRING_INDEX_NONE` is returned if *sub* does not occur. `fn_string_prev_substring()` returns the index of the previous occurrence of *sub* at or before *index* in the string *str*. `FN_STRING_INDEX_NONE` is returned if *sub* does not occur. In both of these functions, *string\_case* specifies whether the search should take case-sensitivity into account.

**ERRORS** `fn_string_str()` returns 0 and sets *status* to `FN_E_INCOMPATIBLE_CODE_SETS` if the given string's representation cannot be converted into the code set of the current locale setting. It is implementation-dependent which code sets can be converted into the code set of the current locale.

Code set mismatches that occur during concatenation, searches, or comparisons are resolved in an implementation-dependent way. When an implementation discovers that arguments to substring searches and comparison operations have incompatible

## FN\_string\_t(3XFN)

code sets, it sets *status* to `FN_E_INCOMPATIBLE_CODE_SETS`. In such cases, `fn_string_from_strings()` returns 0. The returned value for comparison operations when there is code set or locale incompatibility is either negative or positive (greater than 0); it is never 0.

`fn_string_from_contents()` returns 0 and *status* is set to `FN_E_INCOMPATIBLE_CODE_SETS` if the supplied code set and/or locale information are not supported by the XFN implementation.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `xfn(3XFN)`, `attributes(5)`

**NOTES** The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

<b>NAME</b>	getaddrinfo, getnameinfo, freeaddrinfo, gai_strerror – translate between node name and address
<b>SYNOPSIS</b>	<pre>cc [flag ...] file... -lsocket -lnsl [library ...] #include &lt;sys/socket.h&gt; #include &lt;netdb.h&gt;  int <b>getaddrinfo</b>(const char *nodename, const char *servname, const     struct addrinfo *hints, struct addrinfo **res);  int <b>getnameinfo</b>(const struct sockaddr *sa, socklen_t salen, char     *host, size_t hostlen, char *serv, size_t servlen, int flags);  void <b>freeaddrinfo</b>(struct addrinfo *ai);  char *<b>gai_strerror</b>(int errcode);</pre>
<b>DESCRIPTION</b>	<p>These functions perform translations from node name to address and from address to node name in a protocol-independent manner.</p> <p>The <code>getaddrinfo()</code> function performs the node name to address translation. The <code>nodename</code> and <code>servname</code> arguments are pointers to null-terminated strings or NULL. One or both of these arguments must be a non-null pointer. In the normal client scenario, both the <code>nodename</code> and <code>servname</code> are specified. In the normal server scenario, only the <code>servname</code> is specified. A non-null <code>nodename</code> string can be either a node name or a numeric host address string (a dotted-decimal IPv4 address or an IPv6 hex address). A non-null <code>servname</code> string can be either a service name or a decimal port number.</p> <p>The caller can optionally pass an <code>addrinfo</code> structure, pointed to by the third argument, to provide hints concerning the type of socket that the caller supports.</p> <p>The <code>addrinfo</code> structure is defined as:</p> <pre>struct addrinfo { int          ai_flags;          /* AI_PASSIVE, AI_CANONNAME, AI_NUMERICHOST */ int          ai_family;        /* PF_xxx */ int          ai_socktype;      /* SOCK_xxx */ int          ai_protocol;      /* 0 or IPPROTO_xxx for IPv4 and IPv6 */ size_t      ai_addrlen;       /* length of ai_addr */ char        *ai_canonname;     /* canonical name for nodename */ struct sockaddr *ai_addr;      /* binary address */ struct addrinfo *ai_next;     /* next structure in linked list */ };</pre> <p>In this <code>hints</code> structure, all members other than <code>ai_flags</code>, <code>ai_family</code>, <code>ai_socktype</code>, and <code>ai_protocol</code> must be 0 or a null pointer. A value of <code>PF_UNSPEC</code> for <code>ai_family</code> indicates that the caller will accept any protocol family. A value of 0 for <code>ai_socktype</code> indicates that the caller will accept any socket type. A value of 0 for <code>ai_protocol</code> indicates that the caller will accept any protocol. For example, if the caller handles only TCP and not UDP, then the <code>ai_socktype</code> member of the <code>hints</code> structure should be set to <code>SOCK_STREAM</code> when <code>getaddrinfo()</code> is called. If the caller handles only</p>

## getaddrinfo(3SOCKET)

IPv4 and not IPv6, then the `ai_family` member of the `hints` structure should be set to `PF_INET` when `getaddrinfo()` is called. If the third argument to `getaddrinfo()` is a null pointer, it is as if the caller had filled in an `addrinfo` structure initialized to 0 with `ai_family` set to `PF_UNSPEC`.

Upon success, a pointer to a linked list of one or more `addrinfo` structures is returned through the final argument. 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(3SOCKET)` function. 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 `bind(3SOCKET)`. In this case, if the `nodename` argument 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, then the returned socket address structure will be ready for a call to `connect(3SOCKET)` (for a connection-oriented protocol) or either `connect(3SOCKET)`, `sendto(3SOCKET)`, or `sendmsg(3SOCKET)` (for a connectionless protocol). If the `nodename` argument 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, then 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 `nodename`.

If the `AI_NUMERICHOST` bit is set in the `ai_flags` member of the `hints` structure, then a non-null `nodename` string must be a numeric host address string. Otherwise an error of `EAI_NONAME` is returned. This flag prevents any type of name resolution service (such as DNS) from being called.

All of the information returned by `getaddrinfo()` is dynamically allocated: the `addrinfo` structures as well as the socket address structures and canonical node name strings pointed to by the `addrinfo` structures. The `freeaddrinfo()` function is called to return this information to the system the function. For `freeaddrinfo()`, the `addrinfo` structure pointed to by the `ai` argument is freed, along with any dynamic storage pointed to by the structure. This operation is repeated until a null `ai_next` pointer is encountered.

## getaddrinfo(3SOCKET)

To aid applications in printing error messages based on the `EAI_*` codes returned by `getaddrinfo()`, the `gai_strerror()` is defined. The argument is one of the `EAI_*` values defined below and the return value points to a string describing the error. If the argument is not one of the `EAI_*` values, the function still returns a pointer to a string whose contents indicate an unknown error.

The `getnameinfo()` function looks up an IP address and port number provided by the caller in the name service database and system-specific database, and returns text strings for both in buffers provided by the caller. The function indicates successful completion by a 0 return value; a non-zero return value indicates failure.

The first argument, *sa*, 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* argument gives the length of the `sockaddr_in` or `sockaddr_in6` structure.

The function returns the node name associated with the IP address in the buffer pointed to by the *host* argument. The caller provides the size of this buffer with the *hostlen* argument. The service name associated with the port number is returned in the buffer pointed to by *serv*, and the *servlen* argument gives the length of this buffer. The caller specifies not to return either string by providing a 0 value for the *hostlen* or *servlen* arguments. Otherwise, the caller must provide buffers large enough to hold the node name and the service name, including the terminating null characters.

To aid the application in allocating buffers for these two returned strings, the following constants are defined in `<netdb.h>`:

```
#define NI_MAXHOST 1025
#define NI_MAXSERV 32
```

The final argument is a flag that changes the default actions of this function. By default, the fully-qualified domain name (FQDN) for the host is looked up in the name service database and returned. If the flag bit `NI_NOFQDN` is set, only the node name portion of the FQDN is returned for local hosts.

If the flag bit `NI_NUMERICHOST` is set, or if the host's name cannot be located in the name service, the numeric form of the host's address is returned instead of its name, for example, by calling `inet_ntop()` (see `inet(3SOCKET)`) instead of `getipnodebyname(3SOCKET)`. If the flag bit `NI_NAMEREQD` is set, an error is returned if the host's name cannot be located in the name service database.

If the flag bit `NI_NUMERICSERV` is set, the numeric form of the service address is returned (for example, its port number) instead of its name. The two `NI_NUMERIC*` flags are required to support the "-n" flag that many commands provide.

A fifth flag bit, `NI_DGRAM`, specifies that the service is a datagram service, and causes `getservbyport(3SOCKET)` to be called with a second argument of "udp" instead of the default "tcp". This is required for the few ports (for example, 512-514) that have different services for UDP and TCP.

## getaddrinfo(3SOCKET)

	These <code>NI_*</code> flags are defined in <code>&lt;netdb.h&gt;</code> along with the <code>AI_*</code> flags already defined for <code>getaddrinfo()</code> .																						
<b>RETURN VALUES</b>	For <code>getaddrinfo()</code> , if the query is successful, a pointer to a linked list of one or more <code>addrinfo</code> structures is returned by the fourth argument and the function returns 0. If the query fails, a non-zero error code will be returned. For <code>getnameinfo()</code> , if successful, the strings <code>hostname</code> and <code>service</code> are copied into <code>host</code> and <code>serv</code> , respectively. If unsuccessful, zero values for either <code>hostlen</code> or <code>servlen</code> will suppress the associated lookup; in this case no data is copied into the applicable buffer. If <code>gai_strerror()</code> is successful, a pointer to a string containing an error message appropriate for the <code>EAI_*</code> errors is returned. If <code>errcode</code> is not one of the <code>EAI_*</code> values, a pointer to a string indicating an unknown error is returned.																						
<b>ERRORS</b>	The following names are the error values returned by <code>getaddrinfo()</code> and are defined in <code>&lt;netdb.h&gt;</code> :  <table><tr><td><code>EAI_ADDRFAMILY</code></td><td>address family for nodename not supported</td></tr><tr><td><code>EAI_AGAIN</code></td><td>temporary failure in name resolution</td></tr><tr><td><code>EAI_BADFLAGS</code></td><td>invalid value for <code>ai_flags</code></td></tr><tr><td><code>EAI_FAIL</code></td><td>non-recoverable failure in name resolution</td></tr><tr><td><code>EAI_FAMILY</code></td><td><code>ai_family</code> not supported</td></tr><tr><td><code>EAI_MEMORY</code></td><td>memory allocation failure</td></tr><tr><td><code>EAI_NODATA</code></td><td>no address associated with nodename</td></tr><tr><td><code>EAI_NONAME</code></td><td>nodename nor servname provided, or not known</td></tr><tr><td><code>EAI_SERVICE</code></td><td>servname not supported for <code>ai_socktype</code></td></tr><tr><td><code>EAI_SOCKTYPE</code></td><td><code>ai_socktype</code> not supported</td></tr><tr><td><code>EAI_SYSTEM</code></td><td>system error returned in <code>errno</code></td></tr></table>	<code>EAI_ADDRFAMILY</code>	address family for nodename not supported	<code>EAI_AGAIN</code>	temporary failure in name resolution	<code>EAI_BADFLAGS</code>	invalid value for <code>ai_flags</code>	<code>EAI_FAIL</code>	non-recoverable failure in name resolution	<code>EAI_FAMILY</code>	<code>ai_family</code> not supported	<code>EAI_MEMORY</code>	memory allocation failure	<code>EAI_NODATA</code>	no address associated with nodename	<code>EAI_NONAME</code>	nodename nor servname provided, or not known	<code>EAI_SERVICE</code>	servname not supported for <code>ai_socktype</code>	<code>EAI_SOCKTYPE</code>	<code>ai_socktype</code> not supported	<code>EAI_SYSTEM</code>	system error returned in <code>errno</code>
<code>EAI_ADDRFAMILY</code>	address family for nodename not supported																						
<code>EAI_AGAIN</code>	temporary failure in name resolution																						
<code>EAI_BADFLAGS</code>	invalid value for <code>ai_flags</code>																						
<code>EAI_FAIL</code>	non-recoverable failure in name resolution																						
<code>EAI_FAMILY</code>	<code>ai_family</code> not supported																						
<code>EAI_MEMORY</code>	memory allocation failure																						
<code>EAI_NODATA</code>	no address associated with nodename																						
<code>EAI_NONAME</code>	nodename nor servname provided, or not known																						
<code>EAI_SERVICE</code>	servname not supported for <code>ai_socktype</code>																						
<code>EAI_SOCKTYPE</code>	<code>ai_socktype</code> not supported																						
<code>EAI_SYSTEM</code>	system error returned in <code>errno</code>																						
<b>FILES</b>	<code>/etc/inet/hosts</code> <code>/etc/inet/ipnodes</code> <code>/etc/netconfig</code> <code>/etc/nsswitch.conf</code>																						
<b>SEE ALSO</b>	<code>gethostbyname(3NSL)</code> , <code>getipnodebyname(3SOCKET)</code> , <code>htonl(3SOCKET)</code> , <code>inet(3SOCKET)</code> , <code>netdb(3HEAD)</code> , <code>socket(3SOCKET)</code> , <code>hosts(4)</code> , <code>ipnodes(4)</code> , <code>nsswitch.conf(4)</code>																						

<b>NAME</b>	gethostbyname, gethostbyname_r, gethostbyaddr, gethostbyaddr_r, gethostent, gethostent_r, sethostent, endhostent – get network host entry
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;netdb.h&gt;  struct hostent *gethostbyname(const char *name); struct hostent *gethostbyname_r(const char *name, struct hostent     *result, char *buffer, int buflen, int *h_errnop); struct hostent *gethostbyaddr(const char *addr, int len, int type); struct hostent *gethostbyaddr_r(const char *addr, int length, int     type, struct hostent *result, char *buffer, int buflen, int     *h_errnop); struct hostent *gethostent(void); struct hostent *gethostent_r(struct hostent *result, char *buffer,     int buflen, int *h_errnop); int sethostent(int stayopen); int endhostent(void);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to obtain entries describing hosts. An entry may come from any of the sources for hosts specified in the <code>/etc/nsswitch.conf</code> file. See <code>nsswitch.conf(4)</code>. Please take note that these functions have been superseded by the newer functions, <code>getipnodebyname(3SOCKET)</code>, <code>getipnodebyaddr(3SOCKET)</code>, and <code>getaddrinfo(3SOCKET)</code>. The newer functions provide greater portability to applications when multithreading is done or technologies such as IPv6 are used. For example, the functions described below cannot be used with applications targeted to work with IPv6.</p> <p><code>gethostbyname()</code> searches for information for a host with the hostname specified by the character-string parameter <i>name</i>.</p> <p><code>gethostbyaddr()</code> searches for information for a host with a given host address. The parameter <i>type</i> specifies the family of the address. This should be one of the address families defined in <code>&lt;sys/socket.h&gt;</code>. See the NOTES section below for more information. Also see the EXAMPLES section below on how to convert a “.” separated Internet IP address notation into the <i>addr</i> parameter. The parameter <i>len</i> specifies the length of the buffer indicated by <i>addr</i>.</p> <p>All addresses are returned in network order. In order to interpret the addresses, <code>byteorder(3SOCKET)</code> must be used for byte order conversion.</p> <p>The functions <code>sethostent()</code>, <code>gethostent()</code>, and <code>endhostent()</code> are used to enumerate host entries from the database.</p>

## gethostbyname(3NSL)

`sethostent()` sets (or resets) the enumeration to the beginning of the set of host entries. This function should be called before the first call to `gethostent()`. Calls to `gethostbyname()` and `gethostbyaddr()` leave the enumeration position in an indeterminate state. If the *stayopen* flag is non-zero, the system may keep allocated resources such as open file descriptors until a subsequent call to `endhostent()`.

Successive calls to `gethostent()` return either successive entries or `NULL`, indicating the end of the enumeration.

`endhostent()` may be called to indicate that the caller expects to do no further host entry retrieval operations; the system may then deallocate resources it was using. It is still allowed, but possibly less efficient, for the process to call more host retrieval functions after calling `endhostent()`.

### Reentrant Interfaces

The functions `gethostbyname()`, `gethostbyaddr()`, and `gethostent()` use static storage that is reused in each call, making these functions unsafe for use in multi-threaded applications.

The functions `gethostbyname_r()`, `gethostbyaddr_r()`, and `gethostent_r()` provide reentrant interfaces for these operations.

Each reentrant interface performs the same operation as its non-reentrant counterpart, named by removing the “\_r” suffix. The reentrant interfaces, however, use buffers supplied by the caller to store returned results, and are safe for use in both single-threaded and multi-threaded applications.

Each reentrant interface takes the same parameters as its non-reentrant counterpart, as well as the following additional parameters. The parameter *result* must be a pointer to a `struct hostent` structure allocated by the caller. On successful completion, the function returns the host entry in this structure. The parameter *buffer* must be a pointer to a buffer supplied by the caller. This buffer is used as storage space for the host data. All of the pointers within the returned `struct hostent result` point to data stored within this buffer. See RETURN VALUES. The buffer must be large enough to hold all of the data associated with the host entry. The parameter *buflen* should give the size in bytes of the buffer indicated by *buffer*. The parameter *h\_errnop* should be a pointer to an integer. An integer error status value is stored there on certain error conditions. See ERRORS.

For enumeration in multi-threaded applications, the position within the enumeration is a process-wide property shared by all threads. `sethostent()` may be used in a multi-threaded application but resets the enumeration position for all threads. If multiple threads interleave calls to `gethostent_r()`, the threads will enumerate disjoint subsets of the host database.

Like their non-reentrant counterparts, `gethostbyname_r()` and `gethostbyaddr_r()` leave the enumeration position in an indeterminate state.

### RETURN VALUES

Host entries are represented by the `struct hostent` structure defined in `<netdb.h>`:



```

struct hostent {
    char    *h_name;           /* canonical name of host */
    char    **h_aliases;      /* alias list */
    int     h_addrtype;       /* host address type */
    int     h_length;         /* length of address */
    char    **h_addr_list;    /* list of addresses */
};

```

See the **EXAMPLES** section below for information about how to retrieve a “.” separated Internet IP address string from the *h\_addr\_list* field of struct *hostent*.

The functions `gethostbyname()`, `gethostbyname_r()`, `gethostbyaddr()`, and `gethostbyaddr_r()` each return a pointer to a struct *hostent* if they successfully locate the requested entry; otherwise they return `NULL`.

The functions `gethostent()` and `gethostent_r()` each return a pointer to a struct *hostent* if they successfully enumerate an entry; otherwise they return `NULL`, indicating the end of the enumeration.

The functions `gethostbyname()`, `gethostbyaddr()`, and `gethostent()` use static storage, so returned data must be copied before a subsequent call to any of these functions if the data is to be saved.

When the pointer returned by the reentrant functions `gethostbyname_r()`, `gethostbyaddr_r()`, and `gethostent_r()` is not `NULL`, it is always equal to the *result* pointer that was supplied by the caller.

The functions `sethostent()` and `endhostent()` return 0 on success.

## ERRORS

The reentrant functions `gethostbyname_r()`, `gethostbyaddr_r()`, and `gethostent_r()` will return `NULL` and set *errno* to `ERANGE` if the length of the buffer supplied by caller is not large enough to store the result. See `Intro(2)` for the proper usage and interpretation of *errno* in multithreaded applications.

The reentrant functions `gethostbyname_r()` and `gethostbyaddr_r()` set the integer pointed to by *h\_errnop* to one of these values in case of error.

On failures, the non-reentrant functions `gethostbyname()` and `gethostbyaddr()` set a global integer *h\_errno* to indicate one of these error codes (defined in `<netdb.h>`): `HOST_NOT_FOUND`, `TRY_AGAIN`, `NO_RECOVERY`, `NO_DATA`, and `NO_ADDRESS`.

Note however that if a resolver is provided with a malformed address, or if any other error occurs before `gethostbyname()` is resolved, then `gethostbyname()` returns an internal error with a value of `-1`.

`gethostbyname()` will set *h\_errno* to `NETDB_INTERNAL` when it returns a `NULL` value.

## gethostbyname(3NSL)

### EXAMPLES **EXAMPLE 1** Using gethostbyname()

Here is a sample program that gets the canonical name, aliases, and "." separated Internet IP addresses for a given "." separated IP address:

```
#include <stdio.h>
#include <string.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
#include <netdb.h>
main(int argc, const char **argv)
{
    ulong_t addr;
    struct hostent *hp;
    char **p;
    if (argc != 2) {
        (void) printf("usage: %s IP-address\n", argv[0]);
        exit (1);
    }
    if ((int)(addr = inet_addr(argv[1])) == -1) {
        (void) printf("IP-address must be of the form a.b.c.d\n");
        exit (2);
    }
    hp = gethostbyaddr((char *)&addr, sizeof (addr), AF_INET);
    if (hp == NULL) {
        (void) printf("host information for %s not found\n", argv[1]);
        exit (3);
    }
    for (p = hp->h_addr_list; *p != 0; p++) {
        struct in_addr in;
        char **q;
        (void) memcpy(&in.s_addr, *p, sizeof (in.s_addr));
        (void) printf("%s\t%s", inet_ntoa(in), hp->h_name);
        for (q = hp->h_aliases; *q != 0; q++)
            (void) printf(" %s", *q);
        (void) putchar('\n');
    }
    exit (0);
}
```

Note that the above sample program is unsafe for use in multithreaded applications.

**FILES** /etc/hosts  
 /etc/netconfig  
 /etc/nsswitch.conf

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	See "Reentrant Interfaces" in DESCRIPTION.

**SEE ALSO** `Intro(2)`, `Intro(3)`, `byteorder(3SOCKET)`, `inet(3SOCKET)`, `netdir(3NSL)`, `hosts(4)`, `netconfig(4)`, `nsswitch.conf(4)`, `attributes(5)`, `netdb(3HEAD)`

**WARNINGS** The reentrant interfaces `gethostbyname_r()`, `gethostbyaddr_r()`, and `gethostent_r()` are included in this release on an uncommitted basis only, and are subject to change or removal in future minor releases.

**NOTES** Programs that use the interfaces described in this manual page cannot be linked statically since the implementations of these functions employ dynamic loading and linking of shared objects at run time.

In order to ensure that they all return consistent results, `gethostbyname()`, `gethostbyname_r()`, and `netdir_getbyname()` are implemented in terms of the same internal library function. This function obtains the system-wide source lookup policy based on the `inet` family entries in `netconfig(4)` and the `hosts:` entry in `nsswitch.conf(4)`. Similarly, `gethostbyaddr()`, `gethostbyaddr_r()`, and `netdir_getbyaddr()` are implemented in terms of the same internal library function. If the `inet` family entries in `netconfig(4)` have a "-" in the last column for `nametoaddr` libraries, then the entry for `hosts` in `nsswitch.conf` will be used; otherwise the `nametoaddr` libraries in that column will be used, and `nsswitch.conf` will not be consulted.

There is no analogue of `gethostent()` and `gethostent_r()` in the `netdir` functions, so these enumeration functions go straight to the `hosts` entry in `nsswitch.conf`. Thus enumeration may return results from a different source than that used by `gethostbyname()`, `gethostbyname_r()`, `gethostbyaddr()`, and `gethostbyaddr_r()`.

All the functions that return a `struct hostent` must always return the *canonical name* in the `h_name` field. This name, by definition, is the well-known and official hostname shared between all aliases and all addresses. The underlying source that satisfies the request determines the mapping of the input name or address into the set of names and addresses in `hostent`. Different sources might do that in different ways. If there is more than one alias and more than one address in `hostent`, no pairing is implied between them.

## gethostbyname(3NSL)

The system will strive to put the addresses on the same subnet as that of the caller first.

When compiling multi-threaded applications, see *Intro(3), Notes On Multithread Applications*, for information about the use of the `_REENTRANT` flag.

Use of the enumeration interfaces `gethostent()` and `gethostent_r()` is discouraged; enumeration may not be supported for all database sources. The semantics of enumeration are discussed further in `nsswitch.conf(4)`.

The current implementations of these functions only return or accept addresses for the Internet address family (type `AF_INET`).

The form for an address of type `AF_INET` is a `struct in_addr` defined in `<netinet/in.h>`. The functions described in `inet(3SOCKET)`, and illustrated in the `EXAMPLES` section above, are helpful in constructing and manipulating addresses in this form.

<b>NAME</b>	gethostname – get name of current host				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;unistd.h&gt;  int <b>gethostname</b>(char *name, size_t namelen);</pre>				
<b>DESCRIPTION</b>	<p>The <code>gethostname()</code> function returns the standard host name for the current machine. The <code>namelen</code> argument specifies the size of the array pointed to by the <code>name</code> argument. The returned name is null-terminated, except that if <code>namelen</code> is an insufficient length to hold the host name, then the returned name is truncated and it is unspecified whether the returned name is null-terminated.</p> <p>Host names are limited to 255 bytes.</p>				
<b>RETURN VALUES</b>	On successful completion, 0 is returned. Otherwise, -1 is returned.				
<b>ERRORS</b>	No errors are defined.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>uname(1)</code> , <code>gethostid(3C)</code> , <code>attributes(5)</code>				

## getipnodebyname(3SOCKET)

<b>NAME</b>	getipnodebyname, getipnodebyaddr, freehostent – get IP node entry
<b>SYNOPSIS</b>	<pre>cc [flag ...] file ... -lsocket -lnsl [library ...]  #include &lt;sys/socket.h&gt; #include &lt;netdb.h&gt;  struct hostent *getipnodebyname(const char *name, int af, int flags,                                 int *error_num);  struct hostent *getipnodebyaddr(const void *src, size_t len, int af,                                 int *error_num);  void freehostent(struct hostent *ptr);</pre>
<b>DESCRIPTION</b>	<p>The <code>getipnodebyname()</code> function searches the ipnodes database from the beginning and finds the first entry for which the hostname specified by <code>name</code> matches the <code>h_name</code> member. It takes an <code>af</code> argument that specifies the address family, which can be either <code>AF_INET</code> for IPv4 addresses or <code>AF_INET6</code> for IPv6 addresses. The <code>flags</code> argument determines what results will be returned based on the value of <code>flags</code>. If the <code>flags</code> argument is set to 0 (zero), then the default operation of this function is specified as follows:</p> <ul style="list-style-type: none"><li>■ If the <code>af</code> argument is <code>AF_INET</code>, then a query is made for an IPv4 address. If successful, IPv4 addresses are returned and the <code>h_length</code> member of the <code>hostent</code> structure will be 4. Otherwise, the function returns a null pointer.</li><li>■ If the <code>af</code> argument is <code>AF_INET6</code>, then a query is made for an IPv6 address. If successful, IPv6 addresses are returned and the <code>h_length</code> member of the <code>hostent</code> structure will be 16. Otherwise, the function returns a null pointer.</li></ul> <p>The <code>flags</code> argument changes the default actions of the function. You can set the <code>flags</code> argument by logically ORing any of the following values together:</p> <pre>AI_V4MAPPED AI_ALL AI_ADDRCONFIG</pre> <p>Note that a special flags value of <code>AI_DEFAULT</code>, as defined below, should handle most applications. In other words, porting simple applications to use IPv6 replaces the call</p> <pre>hptr = gethostbyname(name);</pre> <p>with</p> <pre>hptr = getipnodebyname(name, AF_INET6, AI_DEFAULT, &amp;error_num);</pre> <p>A <code>flags</code> of 0 implies a strict interpretation of the <code>af</code> argument:</p> <ul style="list-style-type: none"><li>■ If <code>flags</code> is 0 and <code>af</code> is <code>AF_INET</code>, then the caller wants only IPv4 addresses. A query is made for A records. If successful, the IPv4 addresses are returned and the <code>h_length</code> member of the <code>hostent</code> structure will be 4; otherwise, the function returns a null pointer.</li></ul>

- If *flags* is 0, and if *af* is AF\_INET6, then the caller wants only IPv6 addresses. A query is made for AAAA records. If successful, the IPv6 addresses are returned and the *h\_length* member of the *hostent* structure will be 16; otherwise, the function returns a null pointer.

Other constants can be logically-ORed into the *flags* argument, to modify the behavior of the function.

- If the AI\_V4MAPPED flag is specified along with an *af* of AF\_INET6, then the caller can accept IPv4-mapped IPv6 addresses. That is, if no AAAA records are found, then a query is made for A records, and any found are returned as IPv4-mapped IPv6 addresses (*h\_length* is 16). The AI\_V4MAPPED flag is ignored unless *af* equals AF\_INET6.
- The AI\_ALL flag is used in conjunction with the AI\_V4MAPPED flag, and is only used with the IPv6 address family. When AI\_ALL is logically ORed with AI\_V4MAPPED flag then the caller wants all addresses: IPv6 and IPv4-mapped IPv6. A query is first made for AAAA records and if successful, the IPv6 addresses are returned. Another query is then made for A records, and any found are returned as IPv4-mapped IPv6 addresses. *h\_length* is 16. Only if both queries fail does the function return a null pointer. This flag is ignored unless *af* equals AF\_INET6.
- The AI\_ADDRCONFIG flag specifies that a query for AAAA 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. For example, if the node has no IPv6 source addresses configured, and *af* equals AF\_INET6, and the node name being looked up has both AAAA and A records, then:
  1. If only AI\_ADDRCONFIG is specified, the function returns a null pointer.
  2. If AI\_ADDRCONFIG or AI\_V4MAPPED is specified, the A records are returned as IPv4-mapped IPv6 addresses.

The special flags value of AI\_DEFAULT is defined as

```
#define AI_DEFAULT (AI_V4MAPPED | AI_ADDRCONFIG)
```

The `getipnodebyname()` function must allow the *name* argument to be either a node name or a literal address string, that is, a dotted-decimal IPv4 address or an IPv6 hex address. This saves applications from having to call `inet_pton(3SOCKET)` to handle literal address strings.

Four scenarios arise based on the type of literal address string and the value of the *af* argument. The two simple cases are when *name* is a dotted-decimal IPv4 address and *af* equals AF\_INET, or when *name* is an IPv6 hex address and *af* equals AF\_INET6. The members of the returned *hostent* structure are:

<i>h_name</i>	Points to a copy of the name argument
<i>h_aliases</i>	Is a null pointer.

## getipnodebyname(3SOCKET)

	<code>h_addrtype</code>	Is a copy of the <i>af</i> argument.
	<code>h_length</code>	Is either 4 (for <code>AF_INET</code> ) or 16 (for <code>AF_INET6</code> ).
	<code>h_addr_list[0]</code>	Is a pointer to the 4-byte or 16-byte binary address.
	<code>h_addr_list[1]</code>	Is a null pointer
<b>PARAMETERS</b>	<i>af</i>	Address family
	<i>flags</i>	Various flags
	<i>name</i>	Name of host
	<i>error_num</i>	Error storage
	<i>src</i>	Address for lookup
	<i>len</i>	Length of address
	<i>ptr</i>	Pointer to <code>hostent</code> structure

**RETURN VALUES** Upon successful completion, `getipnodebyname()` and `getipnodebyaddr()` return a `hostent` structure. Otherwise they return `NULL`.

The `hostent` structure does not change from its existing definition when used with `gethostbyname(3NSL)`. For example, host entries are represented by the `struct hostent` structure defined in `<netdb.h>`:

```
struct hostent {
    char    *h_name;           /* canonical name of host */
    char    **h_aliases;      /* alias list */
    int     h_addrtype;       /* host address type */
    int     h_length;         /* length of address */
    char    **h_addr_list;    /* list of addresses */
};
```

An error occurs when *name* is an IPv6 hex address and *af* equals `AF_INET`. The function's return value is a null pointer and *error\_num* equals `HOST_NOT_FOUND`.

The `getipnodebyaddr()` function has the same arguments as the existing `gethostbyaddr(3NSL)` function, but adds an error number. As with `getipnodebyname()`, `getipnodebyaddr()` is thread safe. The *error\_num* value is returned to the caller with the appropriate error code to support thread safe error code returns. The following error conditions can be returned for *error\_num*:

<code>HOST_NOT_FOUND</code>	Host is unknown.
<code>NO_DATA</code>	No address is available for the <i>name</i> specified in the server request. This error is not a soft error. Another type of <i>name</i> server request might be successful.
<code>NO_RECOVERY</code>	An unexpected server failure occurred, which is a nonrecoverable error.



TRY\_AGAIN

This error is a soft error that indicates that the local server did not receive a response from an authoritative server. A retry at some later time might be successful.

One possible source of confusion is the handling of IPv4-mapped IPv6 addresses and IPv4-compatible IPv6 addresses, but the following logic should apply:

1. If *af* is AF\_INET6, and if *len* equals 16, and if the IPv6 address is an IPv4-mapped IPv6 address or an IPv4-compatible IPv6 address, then skip over the first 12 bytes of the IPv6 address, set *af* to AF\_INET, and set *len* to 4.
2. If *af* is AF\_INET, lookup the *name* for the given IPv4 address.
3. If *af* is AF\_INET6, lookup the *name* for the given IPv6 address.
4. If the function is returning success, then the single address that is returned in the *hostent* structure is a copy of the first argument to the function with the same address family that was passed as an argument to this function.

All four steps listed are performed in order.

This structure, and the information pointed to by this structure, are dynamically allocated by `getipnodebyname()` and `getipnodebyaddr()`. The `freehostent()` function frees this memory.

## EXAMPLES

**EXAMPLE 1** Getting the Canonical Name, Aliases, and Internet IP Addresses for a Given Hostname

The following is a sample program that retrieves the canonical name, aliases, and all Internet IP addresses, both version 6 and version 4, for a given hostname.

```
#include <stdio.h>
#include <string.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
#include <netdb.h>

main(int argc, const char **argv)
{
    char abuf[INET6_ADDRSTRLEN];
    int error_num;
    struct hostent *hp;
    char **p;

    if (argc != 2) {
        (void) printf("usage: %s hostname\n", argv[0]);
        exit (1);
    }

    /* argv[1] can be a pointer to a hostname or literal IP address */
    hp = getipnodebyname(argv[1], AF_INET6, AI_ALL | AI_ADDRCONFIG |
        AI_V4MAPPED, &error_num);
```

## getipnodebyname(3SOCKET)

**EXAMPLE 1** Getting the Canonical Name, Aliases, and Internet IP Addresses for a Given Hostname *(Continued)*

```

if (hp == NULL) {
    if (error_num == TRY_AGAIN) {
        printf("%s: unknown host or invalid literal address "
              "(try again later)\n", argv[1]);
    } else {
        printf("%s: unknown host or invalid literal address\n",
              argv[1]);
    }
    exit (1);
}
for (p = hp->h_addr_list; *p != 0; p++) {
    struct in6_addr in6;
    char **q;

    bcopy(*p, (caddr_t)&in6, hp->h_length);
    (void) printf("%s\t%s", inet_ntop(AF_INET6, (void *)&in6,
        abuf, sizeof(abuf)), hp->h_name);
    for (q = hp->h_aliases; *q != 0; q++)
        (void) printf(" %s", *q);
    (void) putchar('\n');
}
freehostent(hp);
exit (0);
}

```

**FILES** /etc/inet/hosts  
 /etc/inet/ipnodes  
 /etc/netconfig  
 /etc/nsswitch.conf

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl, SUNWarc (32-bit)
	SUNWcslx (64-bit)
MT Level	Safe

**SEE ALSO** getaddrinfo(3SOCKET), gethostbyname(3NSL), htonl(3SOCKET), inet(3SOCKET), netdb(3HEAD), hosts(4), ipnodes(4), nsswitch.conf(4), attributes(5)

**NOTES** Programs that use the interfaces described in this manual page cannot be linked statically since the implementations of these functions employ dynamic loading and linking of shared objects at run time.

No enumeration functions are provided for IPv6. Existing enumeration functions, for example, `sethostent(3NSL)` does not work in combination with `getipnodebyname()` and `getipnodebyaddr()`.

All the functions that return a `struct hostent` must always return the canonical in the `h_name` field. This name, by definition, is the well-known and official hostname shared between all aliases and all addresses. The underlying source that satisfies the request determines the mapping of the input name or address into the set of names and addresses in `hostent`. Different sources might do that in different ways. If more than one alias and more than one address in `hostent` exist, no pairing is implied between them.

The current implementations of these functions only return or accept addresses for the Internet address family (type `AF_INET`) or the Internet address family Version 6 (type `AF_INET6`).

The form for an address of type `AF_INET` is a `struct in_addr` defined in `<netinet/in.h>`. The form for an address of type `AF_INET6` is a `struct in6_addr`, defined also in `<netinet/in.h>`. The functions described in `inet_ntop(3SOCKET)` and `inet_pton(3SOCKET)` that are illustrated in the **EXAMPLES** section are helpful in constructing and manipulating addresses in either of these forms.

## getnetbyname(3SOCKET)

<b>NAME</b>	getnetbyname, getnetbyname_r, getnetbyaddr, getnetbyaddr_r, getnetent, getnetent_r, setnetent, endnetent – get network entry								
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;netdb.h&gt;  struct netent *getnetbyname(const char *name); struct netent *getnetbyname_r(const char *name, struct netent     *result, char *buffer, int buflen); struct netent *getnetbyaddr(long net, inttype); struct netent *getnetbyaddr_r(long net, inttype, struct netent     *result, char *buffer, int buflen); struct netent *getnetent(void); struct netent *getnetent_r(struct netent *result, char *buffer, int     buflen); int setnetent(int stayopen); int endnetent(void);</pre>								
<b>DESCRIPTION</b>	<p>These functions are used to obtain entries for networks. An entry may come from any of the sources for networks specified in the <code>/etc/nsswitch.conf</code> file. See <code>nsswitch.conf(4)</code>.</p> <p><code>getnetbyname()</code> searches for a network entry with the network name specified by the character string parameter <i>name</i>.</p> <p><code>getnetbyaddr()</code> searches for a network entry with the network address specified by <i>net</i>. The parameter <i>type</i> specifies the family of the address. This should be one of the address families defined in <code>&lt;sys/socket.h&gt;</code>. See the NOTES section below for more information.</p> <p>Network numbers and local address parts are returned as machine format integer values, that is, in host byte order. See also <code>inet_network(3SOCKET)</code>.</p> <p>The <code>netent.n_net</code> member in the <code>netent</code> structure pointed to by the return value of the above functions is calculated by <code>inet_network()</code>. The <code>inet_network()</code> function returns a value in host byte order that is aligned based upon the input string. For example:</p> <table><thead><tr><th>Text</th><th>Value</th></tr></thead><tbody><tr><td>"10"</td><td>0x0000000a</td></tr><tr><td>"10.0"</td><td>0x00000a00</td></tr><tr><td>"10.0.1"</td><td>0a000a0001</td></tr></tbody></table>	Text	Value	"10"	0x0000000a	"10.0"	0x00000a00	"10.0.1"	0a000a0001
Text	Value								
"10"	0x0000000a								
"10.0"	0x00000a00								
"10.0.1"	0a000a0001								

getnetbyname(3SOCKET)

Text	Value
"10.0.1.28"	0x0a000180

Commonly, the alignment of the returned value is used as a crude approximate of pre-CIDR (Classless Inter-Domain Routing) subnet mask. For example:

```
in_addr_t addr, mask;

addr = inet_network(net_name);
mask = ~(in_addr_t)0;
if ((addr & IN_CLASSA_NET) == 0)
    addr <<= 8, mask <<= 8;
if ((addr & IN_CLASSB_NET) == 0)
    addr <<= 8, mask <<= 8;
if ((addr & IN_CLASSC_NET) == 0)
    addr <<= 8, mask <<= 8;
```

This usage is deprecated by the CIDR requirements. See Fuller, V., Li, T., Yu, J., and Varadhan, K. *RFC 1591, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*. Network Working Group. September 1993.

The functions `setnetent()`, `getnetent()`, and `endnetent()` are used to enumerate network entries from the database.

`setnetent()` sets (or resets) the enumeration to the beginning of the set of network entries. This function should be called before the first call to `getnetent()`. Calls to `getnetbyname()` and `getnetbyaddr()` leave the enumeration position in an indeterminate state. If the *stayopen* flag is non-zero, the system may keep allocated resources such as open file descriptors until a subsequent call to `endnetent()`.

Successive calls to `getnetent()` return either successive entries or NULL, indicating the end of the enumeration.

`endnetent()` may be called to indicate that the caller expects to do no further network entry retrieval operations; the system may then deallocate resources it was using. It is still allowed, but possibly less efficient, for the process to call more network entry retrieval functions after calling `endnetent()`.

### Reentrant Interfaces

The functions `getnetbyname()`, `getnetbyaddr()`, and `getnetent()` use static storage that is reused in each call, making these routines unsafe for use in multi-threaded applications.

The functions `getnetbyname_r()`, `getnetbyaddr_r()`, and `getnetent_r()` provide reentrant interfaces for these operations.

Each reentrant interface performs the same operation as its non-reentrant counterpart, named by removing the “\_r” suffix. The reentrant interfaces, however, use buffers supplied by the caller to store returned results, and are safe for use in both single-threaded and multi-threaded applications.

## getnetbyname(3SOCKET)

Each reentrant interface takes the same parameters as its non-reentrant counterpart, as well as the following additional parameters. The parameter *result* must be a pointer to a `struct netent` structure allocated by the caller. On successful completion, the function returns the network entry in this structure. The parameter *buffer* must be a pointer to a buffer supplied by the caller. This buffer is used as storage space for the network entry data. All of the pointers within the returned `struct netent result` point to data stored within this buffer. See RETURN VALUES. The buffer must be large enough to hold all of the data associated with the network entry. The parameter *buflen* should give the size in bytes of the buffer indicated by *buffer*.

For enumeration in multi-threaded applications, the position within the enumeration is a process-wide property shared by all threads. `setnetent()` may be used in a multi-threaded application but resets the enumeration position for all threads. If multiple threads interleave calls to `getnetent_r()`, the threads will enumerate disjointed subsets of the network database.

Like their non-reentrant counterparts, `getnetbyname_r()` and `getnetbyaddr_r()` leave the enumeration position in an indeterminate state.

### RETURN VALUES

Network entries are represented by the `struct netent` structure defined in `<netdb.h>`.

The functions `getnetbyname()`, `getnetbyname_r()`, `getnetbyaddr()`, and `getnetbyaddr_r()` each return a pointer to a `struct netent` if they successfully locate the requested entry; otherwise they return `NULL`.

The functions `getnetent()` and `getnetent_r()` each return a pointer to a `struct netent` if they successfully enumerate an entry; otherwise they return `NULL`, indicating the end of the enumeration.

The functions `getnetbyname()`, `getnetbyaddr()`, and `getnetent()` use static storage, so returned data must be copied before a subsequent call to any of these functions if the data is to be saved.

When the pointer returned by the reentrant functions `getnetbyname_r()`, `getnetbyaddr_r()`, and `getnetent_r()` is non-`NULL`, it is always equal to the *result* pointer that was supplied by the caller.

The functions `setnetent()` and `endnetent()` return 0 on success.

### ERRORS

The reentrant functions `getnetbyname_r()`, `getnetbyaddr_r()` and `getnetent_r()` will return `NULL` and set `errno` to `ERANGE` if the length of the buffer supplied by caller is not large enough to store the result. See `intro(2)` for the proper usage and interpretation of `errno` in multi-threaded applications.

getnetbyname(3SOCKET)

**FILES** /etc/networks  
/etc/nsswitch.conf

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** Intro(2), Intro(3), byteorder(3SOCKET), inet(3SOCKET), netdb(3HEAD), networks(4), nsswitch.conf (4), attributes(5)

Fuller, V., Li, T., Yu, J., and Varadhan, K. *RFC 1591, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*. Network Working Group. September 1993.

**WARNINGS** The reentrant interfaces `getnetbyname_r()`, `getnetbyaddr_r()`, and `getnetent_r()` are included in this release on an uncommitted basis only, and are subject to change or removal in future minor releases.

**NOTES** The current implementation of these functions only return or accept network numbers for the Internet address family (type `AF_INET`). The functions described in `inet(3SOCKET)` may be helpful in constructing and manipulating addresses and network numbers in this form.

Programs that use the interfaces described in this manual page cannot be linked statically since the implementations of these functions employ dynamic loading and linking of shared objects at run time.

When compiling multi-threaded applications, see `Intro(3)`, *Notes On Multithread Applications*, for information about the use of the `_REENTRANT` flag.

Use of the enumeration interfaces `getnetent()` and `getnetent_r()` is discouraged; enumeration may not be supported for all database sources. The semantics of enumeration are discussed further in `nsswitch.conf(4)`.

## getnetconfig(3NSL)

<b>NAME</b>	getnetconfig, setnetconfig, endnetconfig, getnetconfigent, freenetconfigent, nc_perror, nc_sperror – get network configuration database entry
<b>SYNOPSIS</b>	<pre>#include &lt;netconfig.h&gt;  struct netconfig *getnetconfig(void *handlep);  void *setnetconfig(void);  int endnetconfig(void *handlep);  struct netconfig *getnetconfigent(const char *netid);  void freenetconfigent(struct netconfig *netconfigp);  void nc_perror(const char *msg);  char *nc_sperror(void);</pre>
<b>DESCRIPTION</b>	<p>The library routines described on this page are part of the Network Selection component. They provide the application access to the system network configuration database, <code>/etc/netconfig</code>. In addition to the routines for accessing the <code>netconfig</code> database, Network Selection includes the environment variable <code>NETPATH</code> (see <code>environ(5)</code>) and the <code>NETPATH</code> access routines described in <code>getnetpath(3NSL)</code>.</p> <p><code>getnetconfig()</code> returns a pointer to the current entry in the <code>netconfig</code> database, formatted as a <code>struct netconfig</code>. Successive calls will return successive <code>netconfig</code> entries in the <code>netconfig</code> database. <code>getnetconfig()</code> can be used to search the entire <code>netconfig</code> file. <code>getnetconfig()</code> returns <code>NULL</code> at the end of the file. <i>handlep</i> is the handle obtained through <code>setnetconfig()</code>.</p> <p>A call to <code>setnetconfig()</code> has the effect of “binding” to or “rewinding” the <code>netconfig</code> database. <code>setnetconfig()</code> must be called before the first call to <code>getnetconfig()</code> and may be called at any other time. <code>setnetconfig()</code> need <i>not</i> be called before a call to <code>getnetconfigent()</code>. <code>setnetconfig()</code> returns a unique handle to be used by <code>getnetconfig()</code>.</p> <p><code>endnetconfig()</code> should be called when processing is complete to release resources for reuse. <i>handlep</i> is the handle obtained through <code>setnetconfig()</code>. Programmers should be aware, however, that the last call to <code>endnetconfig()</code> frees all memory allocated by <code>getnetconfig()</code> for the <code>struct netconfig</code> data structure. <code>endnetconfig()</code> may not be called before <code>setnetconfig()</code>.</p> <p><code>getnetconfigent()</code> returns a pointer to the <code>struct netconfig</code> structure corresponding to <i>netid</i>. It returns <code>NULL</code> if <i>netid</i> is invalid (that is, does not name an entry in the <code>netconfig</code> database).</p> <p><code>freenetconfigent()</code> frees the <code>netconfig</code> structure pointed to by <i>netconfigp</i> (previously returned by <code>getnetconfigent()</code>).</p> <p><code>nc_perror()</code> prints a message to the standard error indicating why any of the above routines failed. The message is prepended with the string <i>msg</i> and a colon. A <code>NEWLINE</code> is appended at the end of the message.</p>



`nc_spperror()` is similar to `nc_perror()` but instead of sending the message to the standard error, will return a pointer to a string that contains the error message.

`nc_perror()` and `nc_spperror()` can also be used with the `NETPATH` access routines defined in `getnetpath(3NSL)`.

**RETURN VALUES**

`setnetconfig()` returns a unique handle to be used by `getnetconfig()`. In the case of an error, `setnetconfig()` returns `NULL` and `nc_perror()` or `nc_spperror()` can be used to print the reason for failure.

`getnetconfig()` returns a pointer to the current entry in the `netconfig()` database, formatted as a `struct netconfig`. `getnetconfig()` returns `NULL` at the end of the file, or upon failure.

`endnetconfig()` returns 0 on success and -1 on failure (for example, if `setnetconfig()` was not called previously).

On success, `getnetconfig()` returns a pointer to the `struct netconfig` structure corresponding to *netid*; otherwise it returns `NULL`.

`nc_spperror()` returns a pointer to a buffer which contains the error message string. This buffer is overwritten on each call. In multithreaded applications, this buffer is implemented as thread-specific data.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO**

`getnetpath(3NSL)`, `netconfig(4)`, `attributes(5)`, `environ(5)`

## getnetpath(3NSL)

<b>NAME</b>	getnetpath, setnetpath, endnetpath – get /etc/netconfig entry corresponding to NETPATH component
<b>SYNOPSIS</b>	<pre>#include &lt;netconfig.h&gt;  struct netconfig *getnetpath(void *handlep);  void *setnetpath(void);  int endnetpath(void *handlep);</pre>
<b>DESCRIPTION</b>	<p>The routines described on this page are part of the Network Selection component. They provide the application access to the system network configuration database, /etc/netconfig, as it is "filtered" by the NETPATH environment variable. See environ(5). See getnetconfig(3NSL) for other routines that also access the network configuration database directly. The NETPATH variable is a list of colon-separated network identifiers.</p> <p>getnetpath() returns a pointer to the netconfig database entry corresponding to the first valid NETPATH component. The netconfig entry is formatted as a struct netconfig. On each subsequent call, getnetpath() returns a pointer to the netconfig entry that corresponds to the next valid NETPATH component. getnetpath() can thus be used to search the netconfig database for all networks included in the NETPATH variable. When NETPATH has been exhausted, getnetpath() returns NULL.</p> <p>A call to setnetpath() "binds" to or "rewinds" NETPATH. setnetpath() must be called before the first call to getnetpath() and may be called at any other time. It returns a handle that is used by getnetpath().</p> <p>getnetpath() silently ignores invalid NETPATH components. A NETPATH component is invalid if there is no corresponding entry in the netconfig database.</p> <p>If the NETPATH variable is unset, getnetpath() behaves as if NETPATH were set to the sequence of "default" or "visible" networks in the netconfig database, in the order in which they are listed.</p> <p>endnetpath() may be called to "unbind" from NETPATH when processing is complete, releasing resources for reuse. Programmers should be aware, however, that endnetpath() frees all memory allocated by getnetpath() for the struct netconfig data structure. endnetpath() returns 0 on success and -1 on failure (for example, if setnetpath() was not called previously).</p>
<b>RETURN VALUES</b>	<p>setnetpath() returns a handle that is used by getnetpath(). In case of an error, setnetpath() returns NULL. nc_perror() or nc_spperror() can be used to print out the reason for failure. See getnetconfig(3NSL).</p> <p>When first called, getnetpath() returns a pointer to the netconfig database entry corresponding to the first valid NETPATH component. When NETPATH has been exhausted, getnetpath() returns NULL.</p>

getnetpath(3NSL)

endnetpath() returns 0 on success and -1 on failure (for example, if setnetpath() was not called previously).

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** getnetconfig(3NSL), netconfig(4), attributes(5), environ(5)

## getpeername(3SOCKET)

<b>NAME</b>	getpeername – get name of connected peer										
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lsocket -lnsl [ <i>library</i> ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  int <b>getpeername</b>(int <i>s</i>, struct sockaddr *<i>name</i>, socklen_t *<i>namelen</i>);</pre>										
<b>DESCRIPTION</b>	getpeername() returns the name of the peer connected to socket <i>s</i> . The <i>int</i> pointed to by the <i>namelen</i> parameter should be initialized to indicate the amount of space pointed to by <i>name</i> . On return it contains the actual size of the name returned (in bytes), prior to any truncation. The name is truncated if the buffer provided is too small.										
<b>RETURN VALUES</b>	If successful, getpeername() returns 0; otherwise it returns -1 and sets <i>errno</i> to indicate the error.										
<b>ERRORS</b>	The call succeeds unless: <table><tr><td>EBADF</td><td>The argument <i>s</i> is not a valid descriptor.</td></tr><tr><td>ENOMEM</td><td>There was insufficient user memory for the operation to complete.</td></tr><tr><td>ENOSR</td><td>There were insufficient STREAMS resources available for the operation to complete.</td></tr><tr><td>ENOTCONN</td><td>The socket is not connected.</td></tr><tr><td>ENOTSOCK</td><td>The argument <i>s</i> is not a socket.</td></tr></table>	EBADF	The argument <i>s</i> is not a valid descriptor.	ENOMEM	There was insufficient user memory for the operation to complete.	ENOSR	There were insufficient STREAMS resources available for the operation to complete.	ENOTCONN	The socket is not connected.	ENOTSOCK	The argument <i>s</i> is not a socket.
EBADF	The argument <i>s</i> is not a valid descriptor.										
ENOMEM	There was insufficient user memory for the operation to complete.										
ENOSR	There were insufficient STREAMS resources available for the operation to complete.										
ENOTCONN	The socket is not connected.										
ENOTSOCK	The argument <i>s</i> is not a socket.										
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe						
ATTRIBUTE TYPE	ATTRIBUTE VALUE										
MT-Level	Safe										
<b>SEE ALSO</b>	accept(3SOCKET), bind(3SOCKET), getsockname(3SOCKET), socket(3SOCKET), attributes(5), socket(3HEAD)										

<b>NAME</b>	getpeername – get the name of the peer socket																
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;sys/socket.h&gt;  int <b>getpeername</b>(int <i>socket</i>, struct sockaddr *<i>address</i>, socklen_t     *<i>address_len</i>);</pre>																
<b>DESCRIPTION</b>	<p>The <code>getpeername()</code> function retrieves the peer address of the specified socket, stores this address in the <code>sockaddr</code> structure pointed to by the <code>address</code> argument, and stores the length of this address in the object pointed to by the <code>address_len</code> argument.</p> <p>If the actual length of the address is greater than the length of the supplied <code>sockaddr</code> structure, the stored address will be truncated.</p> <p>If the protocol permits connections by unbound clients, and the peer is not bound, then the value stored in the object pointed to by <code>address</code> is unspecified.</p>																
<b>RETURN VALUES</b>	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.																
<b>ERRORS</b>	<p>The <code>getpeername()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EBADF</td> <td>The <code>socket</code> argument is not a valid file descriptor.</td> </tr> <tr> <td>EFAULT</td> <td>The <code>address</code> or <code>address_len</code> parameter can not be accessed or written.</td> </tr> <tr> <td>EINVAL</td> <td>The socket has been shut down.</td> </tr> <tr> <td>ENOTCONN</td> <td>The socket is not connected or otherwise has not had the peer prespecified.</td> </tr> <tr> <td>ENOTSOCK</td> <td>The <code>socket</code> argument does not refer to a socket.</td> </tr> <tr> <td>EOPNOTSUPP</td> <td>The operation is not supported for the socket protocol.</td> </tr> </table> <p>The <code>getpeername()</code> function may fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">ENOBUFS</td> <td>Insufficient resources were available in the system to complete the call.</td> </tr> <tr> <td>ENOSR</td> <td>There were insufficient STREAMS resources available for the operation to complete.</td> </tr> </table>	EBADF	The <code>socket</code> argument is not a valid file descriptor.	EFAULT	The <code>address</code> or <code>address_len</code> parameter can not be accessed or written.	EINVAL	The socket has been shut down.	ENOTCONN	The socket is not connected or otherwise has not had the peer prespecified.	ENOTSOCK	The <code>socket</code> argument does not refer to a socket.	EOPNOTSUPP	The operation is not supported for the socket protocol.	ENOBUFS	Insufficient resources were available in the system to complete the call.	ENOSR	There were insufficient STREAMS resources available for the operation to complete.
EBADF	The <code>socket</code> argument is not a valid file descriptor.																
EFAULT	The <code>address</code> or <code>address_len</code> parameter can not be accessed or written.																
EINVAL	The socket has been shut down.																
ENOTCONN	The socket is not connected or otherwise has not had the peer prespecified.																
ENOTSOCK	The <code>socket</code> argument does not refer to a socket.																
EOPNOTSUPP	The operation is not supported for the socket protocol.																
ENOBUFS	Insufficient resources were available in the system to complete the call.																
ENOSR	There were insufficient STREAMS resources available for the operation to complete.																
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:																
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe												
ATTRIBUTE TYPE	ATTRIBUTE VALUE																
MT-Level	MT-Safe																
<b>SEE ALSO</b>	<code>accept(3XNET)</code> , <code>bind(3XNET)</code> , <code>getsockname(3XNET)</code> , <code>socket(3XNET)</code> , <code>attributes(5)</code>																

## getprotobyname(3SOCKET)

<b>NAME</b>	getprotobyname, getprotobyname_r, getprotobynumber, getprotobynumber_r, getprotoent, getprotoent_r, setprotoent, endprotoent – get protocol entry
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;netdb.h&gt;  struct protoent *getprotobyname(const char *name);  struct protoent *getprotobyname_r(const char *name, struct     protoent *result, char *buffer, int buflen);  struct protoent *getprotobynumber(int proto);  struct protoent *getprotobynumber_r(int proto, struct protoent     *result, char *buffer, int buflen);  struct protoent *getprotoent(void);  struct protoent *getprotoent_r(struct protoent *result, char *buffer,     int buflen);  int setprotoent(int stayopen);  int endprotoent(void);</pre>
<b>DESCRIPTION</b>	<p>These routines return a protocol entry. Two types of interfaces are supported: reentrant (<code>getprotobyname_r()</code>, <code>getprotobynumber_r()</code>, and <code>getprotoent_r()</code>) and non-reentrant (<code>getprotobyname()</code>, <code>getprotobynumber()</code>, and <code>getprotoent()</code>). The reentrant routines may be used in single-threaded applications and are safe for multi-threaded applications, making them the preferred interfaces.</p> <p>The reentrant routines require additional parameters which are used to return results data. <i>result</i> is a pointer to a struct protoent structure and will be where the returned results will be stored. <i>buffer</i> is used as storage space for elements of the returned results. <i>buflen</i> is the size of <i>buffer</i> and should be large enough to contain all returned data. <i>buflen</i> must be at least 1024 bytes.</p> <p><code>getprotobyname_r()</code>, <code>getprotobynumber_r()</code>, and <code>getprotoent_r()</code> each return a protocol entry.</p> <p>The entry may come from one of the following sources: the protocols file (see <code>protocols(4)</code>), the NIS maps “protocols.byname” and “protocols.bynumber”, and the NIS+ table “protocols”. The sources and their lookup order are specified in the <code>/etc/nsswitch.conf</code> file (see <code>nsswitch.conf(4)</code> for details). Some name services such as NIS will return only one name for a host, whereas others such as NIS+ or DNS will return all aliases.</p> <p><code>getprotobyname_r()</code> and <code>getprotobynumber_r()</code> sequentially search from the beginning of the file until a matching protocol name or protocol number is found, or until an EOF is encountered.</p>

## getprotobyname(3SOCKET)

`getprotobyname()` and `getprotobynumber()` have the same functionality as `getprotobyname_r()` and `getprotobynumber_r()` except that a static buffer is used to store returned results. These routines are unsafe in a multi-threaded application.

`getprotoent_r()` enumerates protocol entries: successive calls to `getprotoent_r()` will return either successive protocol entries or NULL. Enumeration may not be supported by some sources. Note that if multiple threads call `getprotoent_r()`, each will retrieve a subset of the protocol database.

`getprotent()` has the same functionality as `getprotent_r()` except that a static buffer is used to store returned results. This routine is unsafe in a multi-threaded application.

`setprotoent()` “rewinds” to the beginning of the enumeration of protocol entries. If the *stayopen* flag is non-zero, resources such as open file descriptors are not deallocated after each call to `getprotobynumber_r()` and `getprotobyname_r()`. Calls to `getprotobyname_r()`, `getprotobyname()`, `getprotobynumber_r()` and `getprotobynumber()` may leave the enumeration in an indeterminate state, so `setprotoent()` should be called before the first `getprotoent_r()` or `getprotoent()`. Note that `setprotoent()` has process-wide scope, and “rewinds” the protocol entries for all threads calling `getprotoent_r()` as well as main-thread calls to `getprotoent()`.

`endprotoent()` may be called to indicate that protocol processing is complete; the system may then close any open protocols file, deallocate storage, and so forth. It is legitimate, but possibly less efficient, to call more protocol routines after `endprotoent()`.

The internal representation of a protocol entry is a `protoent` structure defined in `<netdb.h>` with the following members:

```
char *p_name;
char **p_aliases;
int p_proto;
```

### RETURN VALUES

`getprotobyname_r()`, `getprotobyname()`, `getprotobynumber_r()`, and `getprotobynumber()` return a pointer to a `struct protoent` if they successfully locate the requested entry; otherwise they return NULL.

`getprotoent_r()` and `getprotoent()` return a pointer to a `struct protoent` if they successfully enumerate an entry; otherwise they return NULL, indicating the end of the enumeration.

### ERRORS

`getprotobyname_r()`, `getprotobynumber_r()`, and `getprotoent_r()` will fail if the following is true:

ERANGE           length of the buffer supplied by caller is not large enough to store the result.

## getprotobyname(3SOCKET)

**FILES** /etc/protocols  
/etc/nsswitch.conf

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	See NOTES below.

**SEE ALSO** intro(3), nsswitch.conf(4), protocols(4), attributes(5), netdb(3HEAD)

**NOTES** Although `getprotobyname_r()`, `getprotobyname_r()`, and `getprotoent_r()` are not mentioned by POSIX 1003.1c, they were added to complete the functionality provided by similar thread-safe functions.

When compiling multithreaded applications, see `intro(3)`, *Notes On Multithread Applications*, for information about the use of the `_REENTRANT` flag.

The routines `getprotobyname_r()`, `getprotobyname_r()`, and `getprotoent_r()` are reentrant and multi-thread safe. The reentrant interfaces can be used in single-threaded as well as multi-threaded applications and are therefore the preferred interfaces.

The routines `getprotobyname()`, `getprotobyaddr()`, and `getprotoent()` use static storage, so returned data must be copied if it is to be saved. Because of their use of static storage for returned data, these routines are not safe for multi-threaded applications.

`setprotoent()` and `endprotoent()` have process-wide scope, and are therefore not safe in multi-threaded applications.

Use of `getprotoent_r()` and `getprotoent()` is discouraged; enumeration is well-defined for the protocols file and is supported (albeit inefficiently) for NIS and NIS+, but in general may not be well-defined. The semantics of enumeration are discussed in `nsswitch.conf(4)`.

**BUGS** Only the Internet protocols are currently understood.

Programs that call `getprotobyname_r()` or `getprotobyname_r()` routines cannot be linked statically since the implementation of these routines requires dynamic linker functionality to access shared objects at run time.



<b>NAME</b>	getpublickey, getsecretkey, publickey – retrieve public or secret key				
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpc.h&gt; #include &lt;rpc/key_prot.h&gt;  int <b>getpublickey</b>(const char <i>netname</i>[MAXNETNAMELEN], char   <i>publickey</i>[HEXKEYBYTES+1]);  int <b>getsecretkey</b>(const char <i>netname</i>[MAXNETNAMELEN], char   <i>secretkey</i>[HEXKEYBYTES+1], const char <i>*passwd</i>);</pre>				
<b>DESCRIPTION</b>	<p>getpublickey() and getsecretkey() get public and secret keys for <i>netname</i>. The key may come from one of the following sources:</p> <ul style="list-style-type: none"> <li>■ the /etc/publickey file. See publickey(4).</li> <li>■ the NIS map “publickey.byname” or the NIS+ table “cred.org_dir”. The sources and their lookup order are specified in the /etc/nsswitch.conf file. See nsswitch.conf(4).</li> </ul> <p>getsecretkey() has an extra argument, <i>passwd</i>, which is used to decrypt the encrypted secret key stored in the database.</p>				
<b>RETURN VALUES</b>	Both routines return 1 if they are successful in finding the key. Otherwise, the routines return 0. The keys are returned as null-terminated, hexadecimal strings. If the password supplied to getsecretkey() fails to decrypt the secret key, the routine will return 1 but the <i>secretkey</i> [0] will be set to NULL.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
<b>SEE ALSO</b>	secure_rpc(3NSL), nsswitch.conf(4), publickey(4), attributes(5)				
<b>WARNINGS</b>	If getpublickey() gets the public key from any source other than NIS+, all authenticated NIS+ operations may fail. To ensure that this does not happen, edit the nsswitch.conf(4) file to make sure that the public key is obtained from NIS+.				
<b>NOTES</b>	NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <a href="http://www.sun.com/directory/nisplus/transition.html">http://www.sun.com/directory/nisplus/transition.html</a> .				

## getrpcbyname(3NSL)

<b>NAME</b>	getrpcbyname, getrpcbyname_r, getrpcbynumber, getrpcbynumber_r, getrpcnt, getrpcnt_r, setrpcnt, endrpcnt – get RPC entry
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpc/rpcnt.h&gt;  struct rpcnt *getrpcbyname(const char *name); struct rpcnt *getrpcbyname_r(const char *name, struct rpcnt     *result, char *buffer, int buflen); struct rpcnt *getrpcbynumber(const int number); struct rpcnt *getrpcbynumber_r(const int number, struct rpcnt     *result, char *buffer, int buflen); struct rpcnt *getrpcnt(void); struct rpcnt *getrpcnt_r(struct rpcnt *result, char *buffer, int     buflen); void setrpcnt(const int stayopen); void endrpcnt(void);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to obtain entries for RPC (Remote Procedure Call) services. An entry may come from any of the sources for <code>rpc</code> specified in the <code>/etc/nsswitch.conf</code> file (see <code>nsswitch.conf(4)</code>).</p> <p><code>getrpcbyname()</code> searches for an entry with the RPC service name specified by the parameter <code>name</code>.</p> <p><code>getrpcbynumber()</code> searches for an entry with the RPC program number <code>number</code>.</p> <p>The functions <code>setrpcnt()</code>, <code>getrpcnt()</code>, and <code>endrpcnt()</code> are used to enumerate RPC entries from the database.</p> <p><code>setrpcnt()</code> sets (or resets) the enumeration to the beginning of the set of RPC entries. This function should be called before the first call to <code>getrpcnt()</code>. Calls to <code>getrpcbyname()</code> and <code>getrpcbynumber()</code> leave the enumeration position in an indeterminate state. If the <code>stayopen</code> flag is non-zero, the system may keep allocated resources such as open file descriptors until a subsequent call to <code>endrpcnt()</code>.</p> <p>Successive calls to <code>getrpcnt()</code> return either successive entries or <code>NULL</code>, indicating the end of the enumeration.</p> <p><code>endrpcnt()</code> may be called to indicate that the caller expects to do no further RPC entry retrieval operations; the system may then deallocate resources it was using. It is still allowed, but possibly less efficient, for the process to call more RPC entry retrieval functions after calling <code>endrpcnt()</code>.</p>
<b>Reentrant Interfaces</b>	The functions <code>getrpcbyname()</code> , <code>getrpcbynumber()</code> , and <code>getrpcnt()</code> use static storage that is re-used in each call, making these routines unsafe for use in multithreaded applications.

The functions `getrpcbyname_r()`, `getrpcbynumber_r()`, and `getrpccent_r()` provide reentrant interfaces for these operations.

Each reentrant interface performs the same operation as its non-reentrant counterpart, named by removing the “\_r” suffix. The reentrant interfaces, however, use buffers supplied by the caller to store returned results, and are safe for use in both single-threaded and multithreaded applications.

Each reentrant interface takes the same parameters as its non-reentrant counterpart, as well as the following additional parameters. The parameter *result* must be a pointer to a `struct rpcent` structure allocated by the caller. On successful completion, the function returns the RPC entry in this structure. The parameter *buffer* must be a pointer to a buffer supplied by the caller. This buffer is used as storage space for the RPC entry data. All of the pointers within the returned `struct rpcent result` point to data stored within this buffer (see RETURN VALUES). The buffer must be large enough to hold all of the data associated with the RPC entry. The parameter *buflen* should give the size in bytes of the buffer indicated by *buffer*.

For enumeration in multithreaded applications, the position within the enumeration is a process-wide property shared by all threads. `setrpccent()` may be used in a multithreaded application but resets the enumeration position for all threads. If multiple threads interleave calls to `getrpccent_r()`, the threads will enumerate disjoint subsets of the RPC entry database.

Like their non-reentrant counterparts, `getrpcbyname_r()` and `getrpcbynumber_r()` leave the enumeration position in an indeterminate state.

## RETURN VALUES

RPC entries are represented by the `struct rpcent` structure defined in `<rpc/rpcent.h>`:

```
struct rpcent {
    char *r_name;           /* name of this rpc service
    char **r_aliases;      /* zero-terminated list of alternate names */
    int r_number;          /* rpc program number */
};
```

The functions `getrpcbyname()`, `getrpcbyname_r()`, `getrpcbynumber()`, and `getrpcbynumber_r()` each return a pointer to a `struct rpcent` if they successfully locate the requested entry; otherwise they return `NULL`.

The functions `getrpccent()` and `getrpccent_r()` each return a pointer to a `struct rpcent` if they successfully enumerate an entry; otherwise they return `NULL`, indicating the end of the enumeration.

The functions `getrpcbyname()`, `getrpcbynumber()`, and `getrpccent()` use static storage, so returned data must be copied before a subsequent call to any of these functions if the data is to be saved.

## getrpcbyname(3NSL)

When the pointer returned by the reentrant functions `getrpcbyname_r()`, `getrpcbynumber_r()`, and `getrpccent_r()` is non-NULL, it is always equal to the *result* pointer that was supplied by the caller.

**ERRORS** The reentrant functions `getrpcbyname_r()`, `getrpcbynumber_r()` and `getrpccent_r()` will return NULL and set `errno` to `ERANGE` if the length of the buffer supplied by caller is not large enough to store the result. See `intro(2)` for the proper usage and interpretation of `errno` in multithreaded applications.

**FILES** `/etc/rpc`  
`/etc/nsswitch.conf`

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	See "Reentrant Interfaces" in <code>DESCRIPTION</code> .

**SEE ALSO** `rpcinfo(1M)`, `rpc(3NSL)`, `nsswitch.conf(4)`, `rpc(4)`, `attributes(5)`

**WARNINGS** The reentrant interfaces `getrpcbyname_r()`, `getrpcbynumber_r()`, and `getrpccent_r()` are included in this release on an uncommitted basis only, and are subject to change or removal in future minor releases.

**NOTES** Programs that use the interfaces described in this manual page cannot be linked statically since the implementations of these functions employ dynamic loading and linking of shared objects at run time.

When compiling multithreaded applications, see `intro(3)`, *Notes On Multithreaded Applications*, for information about the use of the `_REENTRANT` flag.

Use of the enumeration interfaces `getrpccent()` and `getrpccent_r()` is discouraged; enumeration may not be supported for all database sources. The semantics of enumeration are discussed further in `nsswitch.conf(4)`.

<b>NAME</b>	getservbyname, getservbyname_r, getservbyport, getservbyport_r, getservent, getservent_r, setservent, endservent – get service entry
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;netdb.h&gt;  struct servent *getservbyname(const char *name, const char *proto); struct servent *getservbyname_r(const char *name, const char     *proto, struct servent *result, char *buffer, int buflen); struct servent *getservbyport(int port, const char *proto); struct servent *getservbyport_r(int port, const char *proto, struct     servent *result, char *buffer, int buflen); struct servent *getservent(void); struct servent *getservent_r(struct servent *result, char *buffer,     int buflen); int setservent(int stayopen); int endservent(void);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to obtain entries for Internet services. An entry may come from any of the sources for services specified in the <code>/etc/nsswitch.conf</code> file. See <code>nsswitch.conf(4)</code>.</p> <p><code>getservbyname()</code> and <code>getservbyport()</code> sequentially search from the beginning of the file until a matching protocol name or port number is found, or until end-of-file is encountered. If a protocol name is also supplied (non-NULL), searches must also match the protocol.</p> <p><code>getservbyname()</code> searches for an entry with the Internet service name specified by the parameter <i>name</i>.</p> <p><code>getservbyport()</code> searches for an entry with the Internet port number <i>port</i>.</p> <p>All addresses are returned in network order. In order to interpret the addresses, <code>byteorder(3SOCKET)</code> must be used for byte order conversion. The string <i>proto</i> is used by both <code>getservbyname()</code> and <code>getservbyport()</code> to restrict the search to entries with the specified protocol. If <i>proto</i> is NULL, entries with any protocol may be returned.</p> <p>The functions <code>setservent()</code>, <code>getservent()</code>, and <code>endservent()</code> are used to enumerate entries from the services database.</p>

## getservbyname(3SOCKET)

`setservent ()` sets (or resets) the enumeration to the beginning of the set of service entries. This function should be called before the first call to `getservent ()`. Calls to the functions `getservbyname ()` and `getservbyport ()` leave the enumeration position in an indeterminate state. If the *stayopen* flag is non-zero, the system may keep allocated resources such as open file descriptors until a subsequent call to `endservent ()`.

`getservent ()` reads the next line of the file, opening the file if necessary. `getservent ()` opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to `getservent ()` (either directly, or indirectly through one of the other "getserv" calls).

Successive calls to `getservent ()` return either successive entries or `NULL`, indicating the end of the enumeration.

`endservent ()` closes the file. `endservent ()` may be called to indicate that the caller expects to do no further service entry retrieval operations; the system may then deallocate resources it was using. It is still allowed, but possibly less efficient, for the process to call more service entry retrieval functions after calling `endservent ()`.

### Reentrant Interfaces

The functions `getservbyname ()`, `getservbyport ()`, and `getservent ()` use static storage that is re-used in each call, making these functions unsafe for use in multithreaded applications.

The functions `getservbyname_r ()`, `getservbyport_r ()`, and `getservent_r ()` provide reentrant interfaces for these operations.

Each reentrant interface performs the same operation as its non-reentrant counterpart, named by removing the "\_r" suffix. The reentrant interfaces, however, use buffers supplied by the caller to store returned results, and are safe for use in both single-threaded and multithreaded applications.

Each reentrant interface takes the same parameters as its non-reentrant counterpart, as well as the following additional parameters. The parameter *result* must be a pointer to a `struct servent` structure allocated by the caller. On successful completion, the function returns the service entry in this structure. The parameter *buffer* must be a pointer to a buffer supplied by the caller. This buffer is used as storage space for the service entry data. All of the pointers within the returned `struct servent result` point to data stored within this buffer. See the RETURN VALUES section of this man page. The buffer must be large enough to hold all of the data associated with the service entry. The parameter *buflen* should give the size in bytes of the buffer indicated by *buffer*.

For enumeration in multithreaded applications, the position within the enumeration is a process-wide property shared by all threads. `setservent ()` may be used in a multithreaded application but resets the enumeration position for all threads. If multiple threads interleave calls to `getservent_r ()`, the threads will enumerate disjoint subsets of the service database.

Like their non-reentrant counterparts, `getservbyname_r()` and `getservbyport_r()` leave the enumeration position in an indeterminate state.

**RETURN VALUES**

Service entries are represented by the `struct servent` structure defined in `<netdb.h>`:

```
struct servent {
    char    *s_name;           /* official name of service */
    char    **s_aliases;      /* alias list */
    int     s_port;           /* port service resides at */
    char    *s_proto;         /* protocol to use */
};
```

The members of this structure are:

`s_name`            The official name of the service.  
`s_aliases`        A zero terminated list of alternate names for the service.  
`s_port`            The port number at which the service resides. Port numbers are returned in network byte order.  
`s_proto`           The name of the protocol to use when contacting the service

The functions `getservbyname()`, `getservbyname_r()`, `getservbyport()`, and `getservbyport_r()` each return a pointer to a `struct servent` if they successfully locate the requested entry; otherwise they return `NULL`.

The functions `getservent()` and `getservent_r()` each return a pointer to a `struct servent` if they successfully enumerate an entry; otherwise they return `NULL`, indicating the end of the enumeration.

The functions `getservbyname()`, `getservbyport()`, and `getservent()` use static storage, so returned data must be copied before a subsequent call to any of these functions if the data is to be saved.

When the pointer returned by the reentrant functions `getservbyname_r()`, `getservbyport_r()`, and `getservent_r()` is non-null, it is always equal to the *result* pointer that was supplied by the caller.

**ERRORS**

The reentrant functions `getservbyname_r()`, `getservbyport_r()` and `getservent_r()` will return `NULL` and set `errno` to `ERANGE` if the length of the buffer supplied by caller is not large enough to store the result. See `intro(2)` for the proper usage and interpretation of `errno` in multithreaded applications.

**FILES**

`/etc/services`            Internet network services  
`/etc/netconfig`           network configuration file  
`/etc/nsswitch.conf`       configuration file for the name-service switch

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

## getservbyname(3SOCKET)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	See "Reentrant Interfaces" in DESCRIPTION.

**SEE ALSO** `intro(2)`, `intro(3)`, `byteorder(3SOCKET)`, `netdir(3NSL)`, `netconfig(4)`, `nsswitch.conf(4)`, `services(4)`, `attributes(5)`, `netdb(3HEAD)`

**WARNINGS** The reentrant interfaces `getservbyname_r()`, `getservbyport_r()`, and `getservent_r()` are included in this release on an uncommitted basis only, and are subject to change or removal in future minor releases.

**NOTES** The functions that return `struct servent` return the least significant 16-bits of the `s_port` field in *network byte order*. `getservbyport()` and `getservbyport_r()` also expect the input parameter `port` in the *network byte order*. See `htons(3SOCKET)` for more details on converting between host and network byte orders.

Programs that use the interfaces described in this manual page cannot be linked statically since the implementations of these functions employ dynamic loading and linking of shared objects at run time.

In order to ensure that they all return consistent results, `getservbyname()`, `getservbyname_r()`, and `netdir_getbyname()` are implemented in terms of the same internal library function. This function obtains the system-wide source lookup policy based on the `inet` family entries in `netconfig(4)` and the `services` entry in `nsswitch.conf(4)`. Similarly, `getservbyport()`, `getservbyport_r()`, and `netdir_getbyaddr()` are implemented in terms of the same internal library function. If the `inet` family entries in `netconfig(4)` have a "-" in the last column for `nametoaddr` libraries, then the entry for `services` in `nsswitch.conf` will be used; otherwise the `nametoaddr` libraries in that column will be used, and `nsswitch.conf` will not be consulted.

There is no analogue of `getservent()` and `getservent_r()` in the `netdir` functions, so these enumeration functions go straight to the `services` entry in `nsswitch.conf`. Thus enumeration may return results from a different source than that used by `getservbyname()`, `getservbyname_r()`, `getservbyport()`, and `getservbyport_r()`.

When compiling multithreaded applications, see `intro(3)`, *Notes On Multithread Applications*, for information about the use of the `_REENTRANT` flag.

Use of the enumeration interfaces `getservent()` and `getservent_r()` is discouraged; enumeration may not be supported for all database sources. The semantics of enumeration are discussed further in `nsswitch.conf(4)`.



**NAME** getsockname – get socket name

**SYNOPSIS**

```
cc [ flag ... ] file ... -lsocket -lnsl [ library ... ]
#include <sys/types.h>
#include <sys/socket.h>

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

**DESCRIPTION** getsockname() returns the current *name* for socket *s*. The *namelen* parameter should be initialized to indicate the amount of space pointed to by *name*. On return it contains the actual size in bytes of the name returned.

**RETURN VALUES** If successful, getsockname() returns 0; otherwise it returns -1 and sets *errno* to indicate the error.

**ERRORS** The call succeeds unless:

EBADF	The argument <i>s</i> is not a valid file descriptor.
ENOMEM	There was insufficient memory available for the operation to complete.
ENOSR	There were insufficient STREAMS resources available for the operation to complete.
ENOTSOCK	The argument <i>s</i> is not a socket.

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** bind(3SOCKET), getpeername(3SOCKET), socket(3SOCKET), attributes(5)

## getsockname(3XNET)

<b>NAME</b>	getsockname – get the socket name														
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  int <b>getsockname</b>(int <i>socket</i>, struct sockaddr *<i>address</i>, socklen_t     *<i>address_len</i>);</pre>														
<b>DESCRIPTION</b>	<p>The <code>getsockname()</code> function retrieves the locally-bound name of the specified socket, stores this address in the <code>sockaddr</code> structure pointed to by the <code>address</code> argument, and stores the length of this address in the object pointed to by the <code>address_len</code> argument.</p> <p>If the actual length of the address is greater than the length of the supplied <code>sockaddr</code> structure, the stored address will be truncated.</p> <p>If the socket has not been bound to a local name, the value stored in the object pointed to by <code>address</code> is unspecified.</p>														
<b>RETURN VALUES</b>	<p>Upon successful completion, 0 is returned, the <code>address</code> argument points to the address of the socket, and the <code>address_len</code> argument points to the length of the address. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.</p>														
<b>ERRORS</b>	<p>The <code>getsockname()</code> function will fail:</p> <table><tr><td>EBADF</td><td>The <code>socket</code> argument is not a valid file descriptor.</td></tr><tr><td>EFAULT</td><td>The <code>address</code> or <code>address_len</code> parameter can not be accessed or written.</td></tr><tr><td>ENOTSOCK</td><td>The <code>socket</code> argument does not refer to a socket.</td></tr><tr><td>EOPNOTSUPP</td><td>The operation is not supported for this socket's protocol.</td></tr></table> <p>The <code>getsockname()</code> function may fail if:</p> <table><tr><td>EINVAL</td><td>The socket has been shut down.</td></tr><tr><td>ENOBUFS</td><td>Insufficient resources were available in the system to complete the call.</td></tr><tr><td>ENOSR</td><td>There were insufficient STREAMS resources available for the operation to complete.</td></tr></table>	EBADF	The <code>socket</code> argument is not a valid file descriptor.	EFAULT	The <code>address</code> or <code>address_len</code> parameter can not be accessed or written.	ENOTSOCK	The <code>socket</code> argument does not refer to a socket.	EOPNOTSUPP	The operation is not supported for this socket's protocol.	EINVAL	The socket has been shut down.	ENOBUFS	Insufficient resources were available in the system to complete the call.	ENOSR	There were insufficient STREAMS resources available for the operation to complete.
EBADF	The <code>socket</code> argument is not a valid file descriptor.														
EFAULT	The <code>address</code> or <code>address_len</code> parameter can not be accessed or written.														
ENOTSOCK	The <code>socket</code> argument does not refer to a socket.														
EOPNOTSUPP	The operation is not supported for this socket's protocol.														
EINVAL	The socket has been shut down.														
ENOBUFS	Insufficient resources were available in the system to complete the call.														
ENOSR	There were insufficient STREAMS resources available for the operation to complete.														
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe										
ATTRIBUTE TYPE	ATTRIBUTE VALUE														
MT-Level	MT-Safe														
<b>SEE ALSO</b>	<code>accept(3XNET)</code> , <code>bind(3XNET)</code> , <code>getpeername(3XNET)</code> , <code>socket(3XNET)</code> , <code>attributes(5)</code>														

<b>NAME</b>	getsockopt, setsockopt – get and set options on sockets										
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  int <b>getsockopt</b>(int s, int level, int optname, void *optval, int *optlen); int <b>setsockopt</b>(int s, int level, int optname, const void *optval, int     optlen);</pre>										
<b>DESCRIPTION</b>	<p>getsockopt () and setsockopt () manipulate options associated with a socket. Options may exist at multiple protocol levels; they are always present at the uppermost “socket” level.</p> <p>When manipulating socket options, the level at which the option resides and the name of the option must be specified. To manipulate options at the “socket” level, <i>level</i> is specified as SOL_SOCKET. To manipulate options at any other level, <i>level</i> is the protocol number of the protocol that controls the option. For example, to indicate that an option is to be interpreted by the TCP protocol, <i>level</i> is set to the TCP protocol number. See getprotobyname(3SOCKET).</p> <p>The parameters <i>optval</i> and <i>optlen</i> are used to access option values for setsockopt (). For getsockopt (), they identify a buffer in which the value(s) for the requested option(s) are to be returned. For getsockopt (), <i>optlen</i> is a value-result parameter, initially containing the size of the buffer pointed to by <i>optval</i>, and modified on return to indicate the actual size of the value returned. Use a 0 <i>optval</i> if no option value is to be supplied or returned.</p> <p><i>optname</i> and any specified options are passed uninterpreted to the appropriate protocol module for interpretation. The include file &lt;&lt;sys/socket.h&gt; contains definitions for the socket-level options described below. Options at other protocol levels vary in format and name.</p> <p>Most socket-level options take an int for <i>optval</i>. For setsockopt (), the <i>optval</i> parameter should be non-zero to enable a boolean option, or zero if the option is to be disabled. SO_LINGER uses a struct linger parameter that specifies the desired state of the option and the linger interval. struct linger is defined in &lt;&lt;sys/socket.h&gt;. struct linger contains the following members:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">l_onoff</td> <td>on = 1/off = 0</td> </tr> <tr> <td>l_linger</td> <td>linger time, in seconds</td> </tr> </table> <p>The following options are recognized at the socket level. Except as noted, each may be examined with getsockopt () and set with setsockopt ().</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">SO_DEBUG</td> <td>enable/disable recording of debugging information</td> </tr> <tr> <td>SO_REUSEADDR</td> <td>enable/disable local address reuse</td> </tr> <tr> <td>SO_KEEPALIVE</td> <td>enable/disable keep connections alive</td> </tr> </table>	l_onoff	on = 1/off = 0	l_linger	linger time, in seconds	SO_DEBUG	enable/disable recording of debugging information	SO_REUSEADDR	enable/disable local address reuse	SO_KEEPALIVE	enable/disable keep connections alive
l_onoff	on = 1/off = 0										
l_linger	linger time, in seconds										
SO_DEBUG	enable/disable recording of debugging information										
SO_REUSEADDR	enable/disable local address reuse										
SO_KEEPALIVE	enable/disable keep connections alive										

## getsockopt(3SOCKET)

SO_DONTROUTE	enable/disable routing bypass for outgoing messages
SO_LINGER	linger on close if data is present
SO_BROADCAST	enable/disable permission to transmit broadcast messages
SO_OOBINLINE	enable/disable reception of out-of-band data in band
SO_SNDBUF	set buffer size for output
SO_RCVBUF	set buffer size for input
SO_DGRAM_ERRIND	application wants delayed error
SO_TYPE	get the type of the socket (get only)
SO_ERROR	get and clear error on the socket (get only)

SO\_DEBUG enables debugging in the underlying protocol modules. SO\_REUSEADDR indicates that the rules used in validating addresses supplied in a `bind(3SOCKET)` call should allow reuse of local addresses. SO\_KEEPAIVE enables the periodic transmission of messages on a connected socket. If the connected party fails to respond to these messages, the connection is considered broken and processes using the socket are notified using a SIGPIPE signal. SO\_DONTROUTE indicates that outgoing messages should bypass the standard routing facilities. Instead, messages are directed to the appropriate network interface according to the network portion of the destination address.

SO\_LINGER controls the action taken when unsent messages are queued on a socket and a `close(2)` is performed. If the socket promises reliable delivery of data and SO\_LINGER is set, the system will block the process on the `close()` attempt until it is able to transmit the data or until it decides it is unable to deliver the information (a timeout period, termed the linger interval, is specified in the `setsockopt()` call when SO\_LINGER is requested). If SO\_LINGER is disabled and a `close()` is issued, the system will process the `close()` in a manner that allows the process to continue as quickly as possible.

The option SO\_BROADCAST requests permission to send broadcast datagrams on the socket. With protocols that support out-of-band data, the SO\_OOBINLINE option requests that out-of-band data be placed in the normal data input queue as received; it will then be accessible with `recv()` or `read()` calls without the MSG\_OOB flag.

SO\_SNDBUF and SO\_RCVBUF are options that adjust the normal buffer sizes allocated for output and input buffers, respectively. The buffer size may be increased for high-volume connections or may be decreased to limit the possible backlog of incoming data. The maximum buffer size for UDP is determined by the value of the `ndd` variable `udp_max_buf`. The maximum buffer size for TCP is determined the value of the `ndd` variable `tcp_max_buf`. Use the `ndd(1M)` utility to determine the current default values. See the *Solaris Tunable Parameters Reference Manual* for information on setting the values of `udp_max_buf` and `tcp_max_buf`.

## getsockopt(3SOCKET)

By default, delayed errors (such as ICMP port unreachable packets) are returned only for connected datagram sockets. `SO_DGRAM_ERRIND` makes it possible to receive errors for datagram sockets that are not connected. When this option is set, certain delayed errors received after completion of a `sendto()` or `sendmsg()` operation will cause a subsequent `sendto()` or `sendmsg()` operation using the same destination address (*to* parameter) to fail with the appropriate error. See `send(3SOCKET)`.

Finally, `SO_TYPE` and `SO_ERROR` are options used only with `getsockopt()`. `SO_TYPE` returns the type of the socket, for example, `SOCK_STREAM`. It is useful for servers that inherit sockets on startup. `SO_ERROR` returns any pending error on the socket and clears the error status. It may be used to check for asynchronous errors on connected datagram sockets or for other asynchronous errors.

**RETURN VALUES** If successful, `getsockopt()` and `setsockopt()` return 0; otherwise, the functions return -1 and set `errno` to indicate the error.

**ERRORS** The `getsockopt()` and `setsockopt()` calls succeed unless:

<code>EBADF</code>	The argument <i>s</i> is not a valid file descriptor.
<code>ENOMEM</code>	There was insufficient memory available for the operation to complete.
<code>ENOPROTOOPT</code>	The option is unknown at the level indicated.
<code>ENOSR</code>	There were insufficient STREAMS resources available for the operation to complete.
<code>ENOTSOCK</code>	The argument <i>s</i> is not a socket.
<code>ENOBUFS</code>	<code>SO_SNDBUF</code> or <code>SO_RCVBUF</code> exceeds a system limit.
<code>EINVAL</code>	Invalid length for <code>IP_OPTIONS</code> .
<code>EHOSTUNREACH</code>	Invalid address for <code>IP_MULTICAST_IF</code> .
<code>EINVAL</code>	Not a multicast address for <code>IP_ADD_MEMBERSHIP</code> and <code>IP_DROP_MEMBERSHIP</code> .
<code>EADDRNOTAVAIL</code>	Bad interface address for <code>IP_ADD_MEMBERSHIP</code> and <code>IP_DROP_MEMBERSHIP</code> .
<code>EADDRINUSE</code>	Address already joined for <code>IP_ADD_MEMBERSHIP</code> .
<code>ENOENT</code>	Address not joined for <code>IP_DROP_MEMBERSHIP</code> .
<code>EPERM</code>	No permissions.
<code>EINVAL</code>	The specified option is invalid at the specified socket level, or the socket has been shut down.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

getsockopt(3SOCKET)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `ndd(1M)`, `close(2)`, `ioctl(2)`, `read(2)`, `bind(3SOCKET)`,  
`getprotobyname(3SOCKET)`, `recv(3SOCKET)`, `send(3SOCKET)`,  
`socket(3SOCKET)`, `attributes(5)`

*Solaris Tunable Parameters Reference Manual*

<b>NAME</b>	getsockopt – get the socket options										
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;sys/socket.h&gt;  int <b>getsockopt</b>(int <i>socket</i>, int <i>level</i>, int <i>option_name</i>, void *<i>option_value</i>,                socklen_t *<i>option_len</i>);</pre>										
<b>DESCRIPTION</b>	<p>The <code>getsockopt()</code> function retrieves the value for the option specified by the <code>option_name</code> argument for the socket specified by the <code>socket</code> argument. If the size of the option value is greater than <code>option_len</code>, the value stored in the object pointed to by the <code>option_value</code> argument will be silently truncated. Otherwise, the object pointed to by the <code>option_len</code> argument will be modified to indicate the actual length of the value.</p> <p>The <code>level</code> argument specifies the protocol level at which the option resides. To retrieve options at the socket level, specify the <code>level</code> argument as <code>SOL_SOCKET</code>. To retrieve options at other levels, supply the appropriate protocol number for the protocol controlling the option. For example, to indicate that an option will be interpreted by the TCP (Transport Control Protocol), set <code>level</code> to the protocol number of TCP, as defined in the <code>&lt;netinet/in.h&gt;</code> header, or as determined by using <code>getprotobyname(3XNET)</code> function.</p> <p>The socket in use may require the process to have appropriate privileges to use the <code>getsockopt()</code> function.</p> <p>The <code>option_name</code> argument specifies a single option to be retrieved. It can be one of the following values defined in <code>&lt;sys/socket.h&gt;</code>:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>SO_DEBUG</code></td> <td>Reports whether debugging information is being recorded. This option stores an <code>int</code> value. This is a boolean option.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>SO_ACCEPTCONN</code></td> <td>Reports whether socket listening is enabled. This option stores an <code>int</code> value.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>SO_BROADCAST</code></td> <td>Reports whether transmission of broadcast messages is supported, if this is supported by the protocol. This option stores an <code>int</code> value. This is a boolean option.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>SO_REUSEADDR</code></td> <td>Reports whether the rules used in validating addresses supplied to <code>bind(3XNET)</code> should allow reuse of local addresses, if this is supported by the protocol. This option stores an <code>int</code> value. This is a boolean option.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>SO_KEEPALIVE</code></td> <td>Reports whether connections are kept active with periodic transmission of messages, if this is supported by the protocol.</td> </tr> </table> <p>If the connected socket fails to respond to these messages, the connection is broken and processes writing to that socket are notified with a <code>SIGPIPE</code> signal. This option stores an <code>int</code> value.</p>	<code>SO_DEBUG</code>	Reports whether debugging information is being recorded. This option stores an <code>int</code> value. This is a boolean option.	<code>SO_ACCEPTCONN</code>	Reports whether socket listening is enabled. This option stores an <code>int</code> value.	<code>SO_BROADCAST</code>	Reports whether transmission of broadcast messages is supported, if this is supported by the protocol. This option stores an <code>int</code> value. This is a boolean option.	<code>SO_REUSEADDR</code>	Reports whether the rules used in validating addresses supplied to <code>bind(3XNET)</code> should allow reuse of local addresses, if this is supported by the protocol. This option stores an <code>int</code> value. This is a boolean option.	<code>SO_KEEPALIVE</code>	Reports whether connections are kept active with periodic transmission of messages, if this is supported by the protocol.
<code>SO_DEBUG</code>	Reports whether debugging information is being recorded. This option stores an <code>int</code> value. This is a boolean option.										
<code>SO_ACCEPTCONN</code>	Reports whether socket listening is enabled. This option stores an <code>int</code> value.										
<code>SO_BROADCAST</code>	Reports whether transmission of broadcast messages is supported, if this is supported by the protocol. This option stores an <code>int</code> value. This is a boolean option.										
<code>SO_REUSEADDR</code>	Reports whether the rules used in validating addresses supplied to <code>bind(3XNET)</code> should allow reuse of local addresses, if this is supported by the protocol. This option stores an <code>int</code> value. This is a boolean option.										
<code>SO_KEEPALIVE</code>	Reports whether connections are kept active with periodic transmission of messages, if this is supported by the protocol.										

## getsockopt(3XNET)

	This is a boolean option.
SO_LINGER	Reports whether the socket lingers on <code>close(2)</code> if data is present. If <code>SO_LINGER</code> is set, the system blocks the process during <code>close(2)</code> until it can transmit the data or until the end of the interval indicated by the <code>l_linger</code> member, whichever comes first. If <code>SO_LINGER</code> is not specified, and <code>close(2)</code> is issued, the system handles the call in a way that allows the process to continue as quickly as possible. This option stores a <code>linger</code> structure.
SO_OOBINLINE	Reports whether the socket leaves received out-of-band data (data marked urgent) in line. This option stores an <code>int</code> value. This is a boolean option.
SO_SNDBUF	Reports send buffer size information. This option stores an <code>int</code> value.
SO_RCVBUF	Reports receive buffer size information. This option stores an <code>int</code> value.
SO_ERROR	Reports information about error status and clears it. This option stores an <code>int</code> value.
SO_TYPE	Reports the socket type. This option stores an <code>int</code> value.
SO_DONTROUTE	Reports whether outgoing messages bypass the standard routing facilities. The destination must be on a directly-connected network, and messages are directed to the appropriate network interface according to the destination address. The effect, if any, of this option depends on what protocol is in use. This option stores an <code>int</code> value. This is a boolean option.

For boolean options, a zero value indicates that the option is disabled and a non-zero value indicates that the option is enabled.

Options at other protocol levels vary in format and name.

The socket in use may require the process to have appropriate privileges to use the `getsockopt()` function.

**RETURN VALUES** Upon successful completion, `getsockopt()` returns 0. Otherwise, `-1` is returned and `errno` is set to indicate the error.

**ERRORS** The `getsockopt()` function will fail if:

`EBADF` The *socket* argument is not a valid file descriptor.



getsockopt(3XNET)

- EFAULT            The *option\_value* or *option\_len* parameter can not be accessed or written.
- EINVAL            The specified option is invalid at the specified socket level.
- ENOPROTOOPT     The option is not supported by the protocol.
- ENOTSOCK         The *socket* argument does not refer to a socket.

The `getsockopt()` function may fail if:

- EACCES            The calling process does not have the appropriate privileges.
- EINVAL            The socket has been shut down.
- ENOBUFS          Insufficient resources are available in the system to complete the call.
- ENOSR             There were insufficient STREAMS resources available for the operation to complete.

**ATTRIBUTES**    See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO**      `close(2)`, `bind(3XNET)`, `endprotoent(3XNET)`, `setsockopt(3XNET)`, `socket(3XNET)`, `attributes`

## gss\_accept\_sec\_context(3GSS)

<b>NAME</b>	<code>gss_accept_sec_context</code> – accept a security context initiated by a peer application
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_accept_sec_context</b>(OM_uint32 *minor_status,     gss_ctx_id_t *context_handle, const gss_cred_id_t     acceptor_cred_handle, const gss_buffer_t input_token, const     gss_channel_bindings_t input_chan_bindings, const gss_name_t     *src_name, gss_OID *mech_type, gss_buffer_t output_token, OM_uint32     *ret_flags, OM_uint32 *time_rec, gss_cred_id_t *delegated_cred_handle);</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_accept_sec_context()</code> function allows a remotely initiated security context between the application and a remote peer to be established. The routine may return an <i>output_token</i>, which should be transferred to the peer application, where the peer application will present it to <code>gss_init_sec_context()</code>. See <code>gss_init_sec_context(3GSS)</code>. If no token need be sent, <code>gss_accept_sec_context()</code> will indicate this by setting the length field of the <i>output_token</i> argument to zero. To complete the context establishment, one or more reply tokens may be required from the peer application; if so, <code>gss_accept_sec_context()</code> will return a status flag of <code>GSS_S_CONTINUE_NEEDED</code>, in which case it should be called again when the reply token is received from the peer application, passing the token to <code>gss_accept_sec_context()</code> by means of the <i>input_token</i> parameters.</p> <p>Portable applications should be constructed to use the token length and return status to determine whether to send or to wait for a token.</p> <p>Whenever <code>gss_accept_sec_context()</code> returns a major status that includes the value <code>GSS_S_CONTINUE_NEEDED</code>, the context is not fully established and the following restrictions apply to the output parameters:</p> <ul style="list-style-type: none"><li>■ The value returned by means of the <i>time_rec</i> parameter is undefined.</li><li>■ Unless the accompanying <i>ret_flags</i> parameter contains the bit <code>GSS_C_PROT_READY_FLAG</code>, which indicates that per-message services may be applied in advance of a successful completion status, the value returned by the <i>mech_type</i> parameter may be undefined until <code>gss_accept_sec_context()</code> returns a major status value of <code>GSS_S_COMPLETE</code>.</li></ul> <p>The values of the <code>GSS_C_DELEG_FLAG</code>, <code>GSS_C_MUTUAL_FLAG</code>, <code>GSS_C_REPLAY_FLAG</code>, <code>GSS_C_SEQUENCE_FLAG</code>, <code>GSS_C_CONF_FLAG</code>, <code>GSS_C_INTEG_FLAG</code> and <code>GSS_C_ANON_FLAG</code> bits returned by means of the <i>ret_flags</i> parameter are values that would be valid if context establishment were to succeed.</p> <p>The values of the <code>GSS_C_PROT_READY_FLAG</code> and <code>GSS_C_TRANS_FLAG</code> bits within <i>ret_flags</i> indicate the actual state at the time <code>gss_accept_sec_context()</code> returns, whether or not the context is fully established. However, applications should not rely</p>

`gss_accept_sec_context(3GSS)`

on this behavior, as `GSS_C_PROT_READY_FLAG` was not defined in Version 1 of the GSS-API. Instead, applications should be prepared to use per-message services after a successful context establishment, based upon the `GSS_C_INTEG_FLAG` and `GSS_C_CONF_FLAG` values.

All other bits within the *ret\_flags* argument are set to zero.

While `gss_accept_sec_context()` returns `GSS_S_CONTINUE_NEEDED`, the values returned by means of the *ret\_flags* argument indicate the services available from the established context. If the initial call of `gss_accept_sec_context()` fails, no context object is created, and the value of the *context\_handle* parameter is set to `GSS_C_NO_CONTEXT`. In the event of a failure on a subsequent call, the security context and the *context\_handle* parameter are left untouched for the application to delete using `gss_delete_sec_context(3GSS)`. During context establishment, the informational status bits `GSS_S_OLD_TOKEN` and `GSS_S_DUPLICATE_TOKEN` indicate fatal errors; GSS-API mechanisms always return them in association with a routine error of `GSS_S_FAILURE`. This pairing requirement did not exist in version 1 of the GSS-API specification, so applications that wish to run over version 1 implementations must special-case these codes.

## PARAMETERS

The parameter descriptions for `gss_accept_sec_context()` follow:

<i>minor_status</i>	The status code returned by the underlying mechanism.
<i>context_handle</i>	The context handle to return to the initiator. This should be set to <code>GSS_C_NO_CONTEXT</code> before the loop begins.
<i>acceptor_cred_handle</i>	The handle for the credentials acquired by the acceptor, typically through <code>gss_acquire_cred()</code> . It may be initialized to <code>GSS_C_NO_CREDENTIAL</code> to indicate a default credential to use. If no default credential is defined, the function returns <code>GSS_C_NO_CRED</code> .
<i>input_token_buffer</i>	Token received from the context initiative.
<i>input_chan_bindings</i>	Optional application-specified bindings. Allows application to securely bind channel identification information to the security context. Set to <code>GSS_C_NO_CHANNEL_BINDINGS</code> if you do not want to use channel bindings.
<i>src_name</i>	The authenticated name of the context initiator. After use, this name should be deallocated by passing it to <code>gss_release_name()</code> . See <code>gss_release_name(3GSS)</code> . If not required, specify <code>NULL</code> .
<i>mech_type</i>	The security mechanism used. Set to <code>NULL</code> if it does not matter which mechanism is used.
<i>output_token</i>	The token to send to the acceptor. Initialize it to <code>GSS_C_NO_BUFFER</code> before the function is called (or its length field set to zero). If the length is zero, no token need be sent.

## gss\_accept\_sec\_context(3GSS)

*ret\_flags*

Contains various independent flags, each of which indicates that the context supports a specific service option. If not needed, specify NULL. Test the returned bit-mask *ret\_flags* value against its symbolic name to determine if the given option is supported by the context. *ret\_flags* may contain one of the following values:

### GSS\_C\_DELEG\_FLAG

If true, delegated credentials are available by means of the *delegated\_cred\_handle* parameter. If false, no credentials were delegated.

### GSS\_C\_MUTUAL\_FLAG

If true, a remote peer asked for mutual authentication. If false, no remote peer asked for mutual authentication.

### GSS\_C\_REPLY\_FLAG

If true, replay of protected messages will be detected. If false, replayed messages will not be detected.

### GSS\_C\_SEQUENCE\_FLAG

If true, out of sequence protected messages will be detected. If false, they will not be detected.

### GSS\_C\_CONF\_FLAG

If true, confidentiality service may be invoked by calling the *gss\_wrap()* routine. If false, no confidentiality service is available by means of *gss\_wrap()*. *gss\_wrap()* will provide message encapsulation, data-origin authentication and integrity services only.

### GSS\_C\_INTEG\_FLAG

If true, integrity service may be invoked by calling either the *gss\_get\_mic(3GSS)* or the *gss\_wrap(3GSS)* routine. If false, per-message integrity service is not available.

### GSS\_C\_ANON\_FLAG

If true, the initiator does not wish to be authenticated. The *src\_name* parameter, if requested, contains an anonymous internal name. If false, the initiator has been authenticated normally.

### GSS\_C\_PROT\_READY\_FLAG

If true, the protection services specified by the states of *GSS\_C\_CONF\_FLAG* and *GSS\_C\_INTEG\_FLAG* are available if the accompanying major status return value is either *GSS\_S\_COMPLETE* or *GSS\_S\_CONTINUE\_NEEDED*. If false, the protection services are available only if the accompanying major status return value is *GSS\_S\_COMPLETE*.

gss\_accept\_sec\_context(3GSS)

GSS\_C\_TRANS\_FLAG

If true, the resultant security context may be transferred to other processes by means of a call to `gss_export_sec_context(3GSS)`. If false, the security context cannot be transferred.

*time\_rec*

The number of sections for which the context will remain value Specify NULL if not required.

*delegated\_cred\_handle*

The credential value for credentials received from the context's initiator. It is valid only if the initiator has requested that the acceptor act as a proxy: that is, if the *ret\_flag* argument resolves to GSS\_C\_DELEG\_FLAG.

## RETURN VALUES

`gss_accept_sec_context()` may return the following status codes:

GSS\_S\_COMPLETE

Successful completion.

GSS\_S\_CONTINUE\_NEEDED

A token from the peer application is required to complete the context, and that `gss_accept_sec_context()` must be called again with that token.

GSS\_S\_DEFECTIVE\_TOKEN

Consistency checks performed on the *input\_token* failed.

GSS\_S\_DEFECTIVE\_CREDENTIAL

Consistency checks performed on the credential failed.

GSS\_S\_NO\_CRED

The supplied credentials were not valid for context acceptance, or the credential handle did not reference any credentials.

GSS\_S\_CREDENTIALS\_EXPIRED

The referenced credentials have expired.

GSS\_S\_BAD\_BINDINGS

The *input\_token* contains different channel bindings than those specified by means of the *input\_chan\_bindings* parameter.

GSS\_S\_NO\_CONTEXT

The supplied context handle did not refer to a valid context.

GSS\_S\_BAD\_SIG

The *input\_token* contains an invalid MIC.

GSS\_S\_OLD\_TOKEN

The *input\_token* was too old. This is a fatal error while establishing context.

GSS\_S\_DUPLICATE\_TOKEN

The *input\_token* is valid, but it is duplicate of a token already processed. This is a fatal error while establishing context.

GSS\_S\_BAD\_MECH

The token received specified a mechanism that is not supported by the implementation or the provided credential.

## gss\_accept\_sec\_context(3GSS)

GSS\_S\_FAILURE

The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the *minor\_status* parameter details the error condition.

### EXAMPLES

#### EXAMPLE 1 Invoking gss\_accept\_sec\_context () Within a Loop

A typical portable caller should always invoke gss\_accept\_sec\_context () within a loop:

```
gss_ctx_id_t context_hdl = GSS_C_NO_CONTEXT;

do {
    receive_token_from_peer(input_token);
    maj_stat = gss_accept_sec_context(&min_stat,
                                     &context_hdl,
                                     cred_hdl,
                                     input_token,
                                     input_bindings,
                                     &client_name,
                                     &mech_type,
                                     output_token,
                                     &ret_flags,
                                     &time_rec,
                                     &deleg_cred);

    if (GSS_ERROR(maj_stat)) {
        report_error(maj_stat, min_stat);
    };
    if (output_token->length != 0) {
        send_token_to_peer(output_token);
        gss_release_buffer(&min_stat, output_token);
    };
    if (GSS_ERROR(maj_stat)) {
        if (context_hdl != GSS_C_NO_CONTEXT)
            gss_delete_sec_context(&min_stat,
                                   &context_hdl,
                                   GSS_C_NO_BUFFER);

        break;
    };
} while (maj_stat & GSS_S_CONTINUE_NEEDED);
```

### ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

`gss_accept_sec_context(3GSS)`

**SEE ALSO** `gss_delete_sec_context(3GSS)`, `gss_export_sec_context(3GSS)`,  
`gss_get_mic(3GSS)`, `gss_init_sec_context(3GSS)`, `gss_release_name(3GSS)`,  
`gss_wrap(3GSS)`, `attributes(5)`

GSS-API Programming Guide

## gss\_acquire\_cred(3GSS)

<b>NAME</b>	<code>gss_acquire_cred</code> – acquire a handle for a pre-existing credential by name				
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_acquire_cred</b>(OM_uint32 *minor_status, const gss_name_t     *desired_name, OM_uint32 time_req, const gss_OID_set desired_mech,     gss_cred_usage_t cred_usage, gss_cred_id_t *output_cred_handle,     gss_OID_set *actual_mechs, OM_uint32 *time_rec);</pre>				
<b>DESCRIPTION</b>	<p>The <code>gss_acquire_cred()</code> function allows an application to acquire a handle for a pre-existing credential by name. This routine is not intended as a function to login to the network; a function for login to the network would involve creating new credentials rather than merely acquiring a handle to existing credentials.</p> <p>If <i>desired_name</i> is <code>GSS_C_NO_NAME</code>, the call is interpreted as a request for a credential handle that will invoke default behavior when passed to <code>gss_init_sec_context(3GSS)</code> (if <i>cred_usage</i> is <code>GSS_C_INITIATE</code> or <code>GSS_C_BOTH</code>) or <code>gss_accept_sec_context(3GSS)</code> (if <i>cred_usage</i> is <code>GSS_C_ACCEPT</code> or <code>GSS_C_BOTH</code>).</p> <p>Normally <code>gss_acquire_cred()</code> returns a credential that is valid only for the mechanisms requested by the <i>desired_mechs</i> argument. However, if multiple mechanisms can share a single credential element, the function returns all the mechanisms for which the credential is valid in the <i>actual_mechs</i> argument.</p> <p><code>gss_acquire_cred()</code> is intended to be used primarily by context acceptors, since the GSS-API routines obtain initiator credentials through the system login process. Accordingly, you may not acquire <code>GSS_C_INITIATE</code> or <code>GSS_C_BOTH</code> credentials by means of <code>gss_acquire_cred()</code> for any name other than <code>GSS_C_NO_NAME</code>. Alternatively, you may acquire <code>GSS_C_INITIATE</code> or <code>GSS_C_BOTH</code> credentials for a name produced when <code>gss_inquire_cred(3GSS)</code> is applied to a valid credential, or when <code>gss_inquire_context(3GSS)</code> is applied to an active context.</p> <p>If credential acquisition is time-consuming for a mechanism, the mechanism may choose to delay the actual acquisition until the credential is required, for example, by <code>gss_init_sec_context(3GSS)</code> or by <code>gss_accept_sec_context(3GSS)</code>. Such mechanism-specific implementations are, however, invisible to the calling application; thus a call of <code>gss_inquire_cred(3GSS)</code> immediately following the call of <code>gss_acquire_cred()</code> will return valid credential data and incur the overhead of a deferred credential acquisition.</p>				
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_acquire_cred()</code> follow:</p> <table><tr><td><i>desired_name</i></td><td>The name of the principal for which a credential should be acquired.</td></tr><tr><td><i>time_req</i></td><td>The number of seconds that credentials remain valid. Specify <code>GSS_C_INDEFINITE</code> to request that the credentials have the maximum permitted lifetime</td></tr></table>	<i>desired_name</i>	The name of the principal for which a credential should be acquired.	<i>time_req</i>	The number of seconds that credentials remain valid. Specify <code>GSS_C_INDEFINITE</code> to request that the credentials have the maximum permitted lifetime
<i>desired_name</i>	The name of the principal for which a credential should be acquired.				
<i>time_req</i>	The number of seconds that credentials remain valid. Specify <code>GSS_C_INDEFINITE</code> to request that the credentials have the maximum permitted lifetime				



gss\_acquire\_cred(3GSS)

<i>desired_mechs</i>	The set of underlying security mechanisms that may be used. GSS_C_NO_OID_SET may be used to obtain a default.
<i>cred_usage</i>	A flag that indicates how this credential should be used. If the flag is GSS_C_ACCEPT, then credentials will be used only to accept security credentials. GSS_C_INITIATE indicates that credentials will be used only to initiate security credentials. If the flag is GSS_C_BOTH, then credentials may be used either to initiate or accept security contexts.
<i>output_cred_handle</i>	The returned credential handle. Resources associated with this credential handle must be released by the application after use with a call to gss_release_cred(3GSS)
<i>actual_mechs</i>	The set of mechanisms for which the credential is valid. Storage associated with the returned OID-set must be released by the application after use with a call to gss_release_oid_set(3GSS). Specify NULL if not required.
<i>time_rec</i>	Actual number of seconds for which the returned credentials will remain valid. Specify NULL if not required.
<i>minor_status</i>	Mechanism specific status code.

**RETURN VALUES**

gss\_acquire\_cred() may return the following status codes:

GSS_S_COMPLETE	Successful completion.
GSS_S_BAD_MECH	An unavailable mechanism has been requested.
GSS_S_BAD_NAME	The value supplied for <i>desired_name</i> parameter is ill formed.
GSS_S_BAD_NAME_TYPE	The type contained within the <i>desired_name</i> parameter is not supported.
GSS_S_CREDENTIALS_EXPIRED	The credentials could not be acquired because they have expired.
GSS_S_NO_CRED	No credentials were found for the specified name.
GSS_S_FAILURE	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.

**ATTRIBUTES**

See attributes(5) for descriptions of the following attributes:

gss\_acquire\_cred(3GSS)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** gss\_accept\_sec\_context(3GSS), gss\_init\_sec\_context(3GSS),  
gss\_inquire\_context(3GSS), gss\_inquire\_cred(3GSS),  
gss\_release\_cred(3GSS), gss\_release\_oid\_set(3GSS), attributes(5)

GSS-API Programming Guide

<b>NAME</b>	gss_add_cred – add a credential-element to a credential
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_add_cred</b>(OM_uint32 *minor_status, const gss_cred_id_t     input_cred_handle, const gss_name_t desired_name, const gss_OID     desired_mech, gss_cred_usage_t cred_usage, OM_uint32 initiator_time_req,     OM_uint32 acceptor_time_req, gss_cred_id_t *output_cred_handle,     gss_OID_set *actual_mechs, OM_uint32 *initiator_time_rec, OM_uint32     *acceptor_time_rec);</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_add_cred()</code> function adds a credential-element to a credential. The credential-element is identified by the name of the principal to which it refers. This routine is not intended as a function to login to the network; a function for login to the network would involve creating new mechanism-specific authentication data rather than merely acquiring a handle to existing data.</p> <p>If the value of <i>desired_name</i> is <code>GSS_C_NO_NAME</code>, the call is interpreted as a request to add a credential element that will invoke default behavior when passed to <code>gss_init_sec_context(3GSS)</code> (if the value of <i>cred_usage</i> is <code>GSS_C_INITIATE</code> or <code>GSS_C_BOTH</code>) or <code>gss_accept_sec_context(3GSS)</code> (if the value of <i>cred_usage</i> is <code>GSS_C_ACCEPT</code> or <code>GSS_C_BOTH</code>).</p> <p>The <code>gss_add_cred()</code> function is expected to be used primarily by context acceptors, since the GSS-API provides mechanism-specific ways to obtain GSS-API initiator credentials through the system login process. Consequently, the GSS-API therefore does not support acquiring <code>GSS_C_INITIATE</code> or <code>GSS_C_BOTH</code> credentials by means of <code>gss_acquire_cred(3GSS)</code> for any name other than <code>GSS_C_NO_NAME</code>, or from name produced by <code>gss_inquire_cred(3GSS)</code> applied to a valid credential or <code>gss_inquire_context(3GSS)</code> applied to an active context.</p> <p>If credential acquisition is time-consuming for a mechanism, the mechanism may choose to delay the actual acquisition until the credential is required, for example, by <code>gss_init_sec_context(3GSS)</code> or by <code>gss_accept_sec_context(3GSS)</code>. Such mechanism-specific implementation decisions are, however, invisible to the calling application; thus a call to <code>gss_inquire_cred(3GSS)</code> immediately following the call of <code>gss_add_cred()</code> will return valid credential data as well as incur the overhead of deferred credential acquisition.</p> <p>The <code>gss_add_cred()</code> routine can be used either to compose a new credential that contains all credential-elements of the original in addition to the newly-acquired credential-element, or to add the new credential-element to an existing credential. If the value of the <i>output_cred_handle</i> parameter argument is <code>NULL</code>, the new credential-element will be added to the credential identified by <i>input_cred_handle</i>; if a valid pointer is specified for the <i>output_cred_handle</i> parameter, a new credential handle will be created.</p>

## gss\_add\_cred(3GSS)

If the value of *input\_cred\_handle* is `GSS_C_NO_CREDENTIAL`, `gss_add_cred()` will compose a credential and set the *output\_cred\_handle* parameter based on the default behavior. That is, the call will have the same effect as if the application had first made a call to `gss_acquire_cred(3GSS)` specifying the same usage and passing `GSS_C_NO_NAME` as the *desired\_name* parameter to obtain an explicit credential handle that incorporates the default behaviors, then passed this credential handle to `gss_add_cred()`, and finally called `gss_release_cred(3GSS)` on the first credential handle.

If the value of the *input\_cred\_handle* parameter is `GSS_C_NO_CREDENTIAL`, you must supply a non-NULL value for the *output\_cred\_handle* parameter.

### PARAMETERS

The parameter descriptions for `gss_acquire_cred()` follow:

<i>minor_status</i>	A mechanism specific status code.
<i>input_cred_handle</i>	The credential to which the credential-element will be added. If <code>GSS_C_NO_CREDENTIAL</code> is specified, the routine will compose the new credential based on default behavior. While the credential-handle is not modified by <code>gss_add_cred()</code> , if <i>output_cred_handle</i> is NULL, the underlying credential will be modified.
<i>desired_name</i>	Name of principal for which a credential should be acquired.
<i>desired_mech</i>	If the value of <i>desired_mech</i> is <code>GSS_C_BOTH</code> , the credential may be used either to initiate or accept security contexts. If the value of <i>desired_mech</i> is <code>GSS_C_INITIATE</code> , the credential will only be used to initiate security contexts. The credential will only be used to accept security contexts, if the value of <i>desired_mech</i> is <code>GSS_C_ACCEPT</code> .
<i>initiator_time_req</i>	The number of seconds that the credential may remain valid for initiating security contexts. This argument is ignored if the composed credentials are of type <code>GSS_C_ACCEPT</code> . Specify <code>GSS_C_INDEFINITE</code> to request that the credentials have the maximum permitted initiator lifetime.
<i>acceptor_time_req</i>	Number of seconds that the credential may remain valid for accepting security contexts. This argument is ignored if the composed credentials are of type <code>GSS_C_INITIATE</code> . Specify <code>GSS_C_INDEFINITE</code> to request that the credentials have the maximum permitted initiator lifetime.
<i>output_cred_handle</i>	The returned credential handle that contains the new credential-element and all the credential-elements from <i>input_cred_handle</i> . If a valid pointer to a <code>gss_cred_id_t</code> is supplied for this parameter, <code>gss_add_cred()</code> creates a new credential handle containing all credential-elements from <i>input_cred_handle</i> and the newly acquired credential-element; if

gss\_add\_cred(3GSS)

NULL is specified for this parameter, the newly acquired credential-element will be added to the credential identified by *input\_cred\_handle*.

The resources associated with any credential handle returned by means of this parameter must be released by the application after use by a call to `gss_release_cred(3GSS)`.

*actual\_mechs* The complete set of mechanisms for which the new credential is valid. Storage for the returned OID-set must be freed by the application after use by a call to `gss_release_oid_set(3GSS)`. Specify NULL if this parameter is not required.

*initiator\_time\_rec* The actual number of seconds for which the returned credentials will remain valid for initiating contexts using the specified mechanism. If a mechanism does not support expiration of credentials, the value `GSS_C_INDEFINITE` will be returned. Specify NULL if this parameter is not required.

*acceptor\_time\_rec* The actual number of seconds for which the returned credentials will remain valid for accepting security contexts using the specified mechanism. If a mechanism does not support expiration of credentials, the value `GSS_C_INDEFINITE` will be returned. Specify NULL if this parameter is not required.

## RETURN VALUES

`gss_acquire_cred()` may return the following status codes:

<code>GSS_S_COMPLETE</code>	Successful completion.
<code>GSS_S_BAD_MECH</code>	An unavailable mechanism has been requested.
<code>GSS_S_BAD_NAME</code>	The type contained within the <i>desired_name</i> parameter is not supported.
<code>GSS_S_BAD_NAME_TYPE</code>	The value supplied for <i>desired_name</i> parameter is ill formed.
<code>GSS_S_DUPLICATE_ELEMENT</code>	The credential already contains an element for the requested mechanism that has overlapping usage and validity period.
<code>GSS_S_CREDENTIALS_EXPIRED</code>	The credentials could not be added because they have expired.
<code>GSS_S_NO_CRED</code>	No credentials were found for the specified name.
<code>GSS_S_FAILURE</code>	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.

gss\_add\_cred(3GSS)

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_accept_sec_context(3GSS)`, `gss_acquire_cred(3GSS)`,  
`gss_init_sec_context(3GSS)`,  
`gss_inquire_context(3GSS)`, `gss_inquire_cred(3GSS)`,  
`gss_release_cred(3GSS)`, `gss_release_oid_set(3GSS)`, `attributes(5)`

*GSS-API Programming Guide*

<b>NAME</b>	gss_add_oid_set_member – add an object identifier to an object identifier set								
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_add_oid_set_member(OM_uint32 *minor_status, const     gss_OID member_oid, gss_OID_set *oid_set);</pre>								
<b>DESCRIPTION</b>	<p>The <code>gss_add_oid_set_member()</code> function adds an object identifier to an object identifier set. You should use this function in conjunction with <code>gss_create_empty_oid_set(3GSS)</code> when constructing a set of mechanism OIDs for input to <code>gss_acquire_cred(3GSS)</code>. The <code>oid_set</code> parameter must refer to an OID-set created by GSS-API, that is, a set returned by <code>gss_create_empty_oid_set(3GSS)</code>.</p> <p>The GSS-API creates a copy of the <code>member_oid</code> and inserts this copy into the set, expanding the storage allocated to the OID-set elements array, if necessary. The function may add the new member OID anywhere within the elements array, and the GSS-API verifies that the new <code>member_oid</code> is not already contained within the elements array. If the <code>member_oid</code> is already present, the <code>oid_set</code> should remain unchanged.</p>								
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_add_oid_set_member()</code> follow:</p> <p><code>minor_status</code>      A mechanism specific status code.</p> <p><code>member_oid</code>        Object identifier to be copied into the set.</p> <p><code>oid_set</code>            Set in which the object identifier should be inserted.</p>								
<b>RETURN VALUES</b>	<p>The <code>gss_add_oid_set_member()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>    Successful completion.</p> <p><code>GSS_S_FAILURE</code>    The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <code>minor_status</code> parameter details the error condition.</p>								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWgss (32-bit)</td> </tr> <tr> <td></td> <td>SUNWgssx (64-bit)</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<p><code>gss_acquire_cred(3GSS)</code>, <code>gss_create_empty_oid_set(3GSS)</code>, <code>attributes(5)</code></p> <p>GSS-API Programming Guide</p>								

## gss\_canonicalize\_name(3GSS)

<b>NAME</b>	gss_canonicalize_name – convert an internal name to a mechanism name				
<b>SYNOPSIS</b>	<pre>cc [flag...] file... -lgss [library...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_canonicalize_name(OM_uint32 *minor_status, const     gss_name_t input_name, const gss_OID mech_type, gss_name_t     *output_name) ;</pre>				
<b>DESCRIPTION</b>	<p>The <code>gss_canonicalize_name()</code> function generates a canonical mechanism name from an arbitrary internal name. The mechanism name is the name that would be returned to a context acceptor on successful authentication of a context where the initiator used the <code>input_name</code> in a successful call to <code>gss_acquire_cred(3GSS)</code>, specifying an OID set containing <code>mech_type</code> as its only member, followed by a call to <code>gss_init_sec_context(3GSS)</code>, specifying <code>mech_type</code> as the authentication mechanism.</p>				
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_canonicalize_name()</code> follow:</p> <p><i>minor_status</i>            Mechanism-specific status code.</p> <p><i>input_name</i>              The name for which a canonical form is desired.</p> <p><i>mech_type</i>                The authentication mechanism for which the canonical form of the name is desired. The desired mechanism must be specified explicitly; no default is provided.</p> <p><i>output_name</i>              The resultant canonical name. Storage associated with this name must be freed by the application after use with a call to <code>gss_release_name(3GSS)</code>.</p>				
<b>RETURN VALUES</b>	<p>The <code>gss_canonicalize_name()</code> function may return the status codes:</p> <p>GSS_S_COMPLETE            Successful completion.</p> <p>GSS_S_BAD_MECH            The identified mechanism is not supported.</p> <p>GSS_S_BAD_NAME            The provided internal name contains no elements that could be processed by the specified mechanism.</p> <p>GSS_S_BAD_NAME            The provided internal name was ill-formed.</p> <p>GSS_S_FAILURE             The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <code>minor_status</code> parameter details the error condition.</p>				
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWgss (32-bit)</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWgss (32-bit)				



gss\_canonicalize\_name(3GSS)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_acquire_cred(3GSS)`, `gss_init_sec_context(3GSS)`,  
`gss_release_name(3GSS)`, `attributes(5)`

GSS-API Programming Guide

## gss\_compare\_name(3GSS)

<b>NAME</b>	gss_compare_name – compare two internal-form names								
<b>SYNOPSIS</b>	<pre>cc [flag...] file... -lgss [library...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_compare_name(OM_uint32 *minor_status, const gss_name_t     name1, const gss_name_t name2, int *name_equal);</pre>								
<b>DESCRIPTION</b>	<p>The <code>gss_compare_name()</code> function allows an application to compare two internal-form names to determine whether they refer to the same entity.</p> <p>If either name presented to <code>gss_compare_name()</code> denotes an anonymous principal, the routines indicate that the two names do not refer to the same identity.</p>								
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_compare_name()</code> follow:</p> <table><tr><td><i>minor_status</i></td><td>Mechanism-specific status code.</td></tr><tr><td><i>name1</i></td><td>Internal-form name.</td></tr><tr><td><i>name2</i></td><td>Internal-form name.</td></tr><tr><td><i>name_equal</i></td><td>If non-zero, the names refer to same entity. If 0, the names refer to different entities. Strictly, the names are not known to refer to the same identity.</td></tr></table>	<i>minor_status</i>	Mechanism-specific status code.	<i>name1</i>	Internal-form name.	<i>name2</i>	Internal-form name.	<i>name_equal</i>	If non-zero, the names refer to same entity. If 0, the names refer to different entities. Strictly, the names are not known to refer to the same identity.
<i>minor_status</i>	Mechanism-specific status code.								
<i>name1</i>	Internal-form name.								
<i>name2</i>	Internal-form name.								
<i>name_equal</i>	If non-zero, the names refer to same entity. If 0, the names refer to different entities. Strictly, the names are not known to refer to the same identity.								
<b>RETURN VALUES</b>	<p>The <code>gss_compare_name()</code> function may return the following status codes:</p> <table><tr><td>GSS_S_COMPLETE</td><td>Successful completion.</td></tr><tr><td>GSS_S_BAD_NAME_TYPE</td><td>The two names were of incomparable types.</td></tr><tr><td>GSS_S_BAD_NAME</td><td>One or both of <i>name1</i> or <i>name2</i> was ill-formed.</td></tr><tr><td>GSS_S_FAILURE</td><td>The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</td></tr></table>	GSS_S_COMPLETE	Successful completion.	GSS_S_BAD_NAME_TYPE	The two names were of incomparable types.	GSS_S_BAD_NAME	One or both of <i>name1</i> or <i>name2</i> was ill-formed.	GSS_S_FAILURE	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.
GSS_S_COMPLETE	Successful completion.								
GSS_S_BAD_NAME_TYPE	The two names were of incomparable types.								
GSS_S_BAD_NAME	One or both of <i>name1</i> or <i>name2</i> was ill-formed.								
GSS_S_FAILURE	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWgss (32-bit)</td></tr><tr><td></td><td>SUNWgssx (64-bit)</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<p><code>attributes(5)</code> GSS-API Programming Guide</p>								

<b>NAME</b>	gss_context_time – determine how long a context will remain valid								
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_context_time(OM_uint32 *minor_status, gss_ctx_id_t     *context_handle, OM_uint32 *time_rec);</pre>								
<b>DESCRIPTION</b>	The gss_context_time() function determines the number of seconds for which the specified context will remain valid.								
<b>PARAMETERS</b>	<p>The parameter descriptions for gss_context_time() are as follows:</p> <p><i>minor_status</i>      A mechanism-specific status code.</p> <p><i>context_handle</i>      A read-only value. Identifies the context to be interrogated.</p> <p><i>time_rec</i>              Modifies the number of seconds that the context remains valid. If the context has already expired, returns zero.</p>								
<b>RETURN VALUES</b>	<p>The gss_context_time() function returns one of the following status codes:</p> <p>GSS_S_COMPLETE              Successful completion.</p> <p>GSS_S_CONTEXT_EXPIRED      The context has already expired.</p> <p>GSS_S_NO_CONTEXT            The <i>context_handle</i> parameter did not identify a valid context.</p> <p>GSS_S_FAILURE                The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>								
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:								
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWgss (32-bit)</td> </tr> <tr> <td></td> <td>SUNWgssx (64-bit)</td> </tr> <tr> <td>MT Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT Level	Safe								
<b>SEE ALSO</b>	<p>gss_init_sec_context(3GSS), gss_accept_sec_context(3GSS),  gss_delete_sec_context(3GSS), gss_process_context_token(3GSS),  gss_inquire_context(3GSS), gss_wrap_size_limit(3GSS),  gss_export_sec_context(3GSS), gss_import_sec_context(3GSS),  attributes(5)</p> <p>GSS-API Programming Guide</p>								

## gss\_create\_empty\_oid\_set(3GSS)

<b>NAME</b>	<code>gss_create_empty_oid_set</code> – create an object-identifier set containing no object identifiers								
<b>SYNOPSIS</b>	<pre>cc -flag ... file ...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_create_empty_oid_set</b>(OM_uint32 *minor_status,     gss_OID_set *oid_set);</pre>								
<b>DESCRIPTION</b>	The <code>gss_create_empty_oid_set()</code> function creates an object-identifier set containing no object identifiers to which members may be subsequently added using the <code>gss_add_oid_set_member(3GSS)</code> function. These functions can be used to construct sets of mechanism object identifiers for input to <code>gss_acquire_cred(3GSS)</code> .								
<b>PARAMETERS</b>	The parameter descriptions for <code>gss_create_empty_oid_set()</code> follow: <table><tr><td><i>minor_status</i></td><td>Mechanism-specific status code</td></tr><tr><td><i>oid_set</i></td><td>Empty object identifier set. The function will allocate the <code>gss_OID_set_desc</code> object, which the application must free after use with a call to <code>gss_release_oid_set(3GSS)</code>.</td></tr></table>	<i>minor_status</i>	Mechanism-specific status code	<i>oid_set</i>	Empty object identifier set. The function will allocate the <code>gss_OID_set_desc</code> object, which the application must free after use with a call to <code>gss_release_oid_set(3GSS)</code> .				
<i>minor_status</i>	Mechanism-specific status code								
<i>oid_set</i>	Empty object identifier set. The function will allocate the <code>gss_OID_set_desc</code> object, which the application must free after use with a call to <code>gss_release_oid_set(3GSS)</code> .								
<b>RETURN VALUES</b>	The <code>gss_create_empty_oid_set()</code> function may return the following status codes: <table><tr><td><code>GSS_S_COMPLETE</code></td><td>Successful completion</td></tr><tr><td><code>GSS_S_FAILURE</code></td><td>The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</td></tr></table>	<code>GSS_S_COMPLETE</code>	Successful completion	<code>GSS_S_FAILURE</code>	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.				
<code>GSS_S_COMPLETE</code>	Successful completion								
<code>GSS_S_FAILURE</code>	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.								
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWgss (32-bit)</td></tr><tr><td></td><td>SUNWgssx (64-bit)</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<code>gss_acquire_cred(3GSS)</code> , <code>gss_add_oid_set_member(3GSS)</code> , <code>gss_release_oid_set(3GSS)</code> , <code>attributes(5)</code>  GSS-API Programming Guide								

<b>NAME</b>	gss_delete_sec_context – delete a GSS-API security context
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_delete_sec_context(OM_uint32 *minor_status,     gss_ctx_id_t *context_handle, gss_buffer_t output_token);</pre>
<b>DESCRIPTION</b>	<p>Use the <code>gss_delete_sec_context()</code> function to delete a security context. The <code>gss_delete_sec_context()</code> function will delete the local data structures associated with the specified security context. You may not obtain further security services that use the context specified by <code>context_handle</code>.</p> <p>In addition to deleting established security contexts, <code>gss_delete_sec_context()</code> will delete any half-built security contexts that result from incomplete sequences of calls to <code>gss_init_sec_context(3GSS)</code> and <code>gss_accept_sec_context(3GSS)</code>.</p> <p>The Solaris implementation of the GSS-API retains the <code>output_token</code> parameter for compatibility with version 1 of the GSS-API. Both peer applications should invoke <code>gss_delete_sec_context()</code>, passing the value <code>GSS_C_NO_BUFFER</code> to the <code>output_token</code> parameter; this indicates that no token is required. If the application passes a valid buffer to <code>gss_delete_sec_context()</code>, it will return a zero-length token, indicating that no token should be transferred by the application.</p>
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_delete_sec_context()</code> follow:</p> <p><i>minor_status</i>      A mechanism specific status code.</p> <p><i>context_handle</i>      Context handle identifying specific context to delete. After deleting the context, the GSS-API will set <code>context_handle</code> to <code>GSS_C_NO_CONTEXT</code>.</p> <p><i>output_token</i>      A token to be sent to remote applications that instructs them to delete the context.</p>
<b>RETURN VALUES</b>	<p><code>gss_delete_sec_context()</code> may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>      Successful completion.</p> <p><code>GSS_S_NO_CONTEXT</code>      No valid context was supplied.</p> <p><code>GSS_S_FAILURE</code>      The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <code>minor_status</code> parameter details the error condition.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)

`gss_delete_sec_context(3GSS)`

ATTRIBUTE TYPE	ATTRIBUTE VALUE
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_accept_sec_context(3GSS)`, `gss_init_sec_context(3GSS)`, `attributes(5)`

GSS-API Programming Guide

<b>NAME</b>	gss_display_name – convert internal-form name to text
<b>SYNOPSIS</b>	<pre>cc [flag...] file... -lgss [library...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_display_name(OM_uint32 *minor_status, const gss_name_t     input_name, gss_buffer_t output_name_buffer, gss_OID     *output_name_type) ;</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_display_name()</code> function allows an application to obtain a textual representation of an opaque internal-form name for display purposes.</p> <p>If <code>input_name</code> denotes an anonymous principal, the GSS-API returns the <code>gss_OID</code> value <code>GSS_C_NT_ANONYMOUS</code> as the <code>output_name_type</code>, and a textual name that is syntactically distinct from all valid supported printable names in <code>output_name_buffer</code>.</p> <p>If <code>input_name</code> was created by a call to <code>gss_import_name(3GSS)</code>, specifying <code>GSS_C_NO_OID</code> as the name-type, the GSS-API returns <code>GSS_C_NO_OID</code> by means of the <code>output_name_type</code> parameter.</p>
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_display_name()</code> follow:</p> <p><i>minor_status</i>            Mechanism-specific status code.</p> <p><i>input_name</i>             Name in internal form.</p> <p><i>output_name_buffer</i>     Buffer to receive textual name string. The application must free storage associated with this name after use with a call to <code>gss_release_buffer(3GSS)</code>.</p> <p><i>output_name_type</i>       The type of the returned name. The returned <code>gss_OID</code> will be a pointer into static storage and should be treated as read-only by the caller. In particular, the application should not attempt to free it. Specify <code>NULL</code> if this parameter is not required.</p>
<b>RETURN VALUES</b>	<p>The <code>gss_display_name()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>        Successful completion.</p> <p><code>GSS_S_BAD_NAME</code>        The <code>input_name</code> was ill-formed.</p> <p><code>GSS_S_FAILURE</code>        The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <code>minor_status</code> parameter details the error condition.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)

`gss_display_name(3GSS)`

ATTRIBUTE TYPE	ATTRIBUTE VALUE
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_import_name(3GSS)`, `gss_release_buffer(3GSS)`, `attributes(5)`

GSS-API Programming Guide



<b>NAME</b>	gss_display_status – convert a GSS-API status code to text												
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_display_status(OM_uint32 *minor_status, OM_uint32     status_value, int status_type, const gss_OID mech_type, OM_uint32     *message_context, gss_buffer_t status_string);</pre>												
<b>DESCRIPTION</b>	<p>The <code>gss_display_status()</code> function enables an application to obtain a textual representation of a GSS-API status code for display to the user or for logging purposes. Because some status values may indicate multiple conditions, applications may need to call <code>gss_display_status()</code> multiple times, with each call generating a single text string.</p> <p>The <code>message_context</code> parameter is used by <code>gss_acquire_cred()</code> to store state information on error messages that are extracted from a given <code>status_value</code>. The <code>message_context</code> parameter must be initialized to 0 by the application prior to the first call, and <code>gss_display_status()</code> will return a non-zero value in this parameter if there are further messages to extract.</p> <p>The <code>message_context</code> parameter contains all state information required by <code>gss_display_status()</code> to extract further messages from the <code>status_value</code>. If a non-zero value is returned in this parameter, the application is not required to call <code>gss_display_status()</code> again unless subsequent messages are desired.</p>												
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_display_status()</code> follow:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><i>minor_status</i></td> <td>Status code returned by the underlying mechanism.</td> </tr> <tr> <td><i>status_value</i></td> <td>Status value to be converted.</td> </tr> <tr> <td><i>status_type</i></td> <td>If the value is <code>GSS_C_GSS_CODE</code>, <i>status_value</i> is a GSS-API status code. If the value is <code>GSS_C_MECH_CODE</code>, then <i>status_value</i> is a mechanism status code.</td> </tr> <tr> <td><i>mech_type</i></td> <td>Underlying mechanism that is used to interpret a minor status value. Supply <code>GSS_C_NO_OID</code> to obtain the system default.</td> </tr> <tr> <td><i>message_context</i></td> <td>Should be initialized to zero prior to the first call. On return from <code>gss_display_status()</code>, a non-zero <i>status_value</i> parameter indicates that additional messages may be extracted from the status code by means of subsequent calls to <code>gss_display_status()</code>, passing the same <i>status_value</i>, <i>status_type</i>, <i>mech_type</i>, and <i>message_context</i> parameters.</td> </tr> <tr> <td><i>status_string</i></td> <td>Textual representation of the <i>status_value</i>. Storage associated with this parameter must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code>.</td> </tr> </table>	<i>minor_status</i>	Status code returned by the underlying mechanism.	<i>status_value</i>	Status value to be converted.	<i>status_type</i>	If the value is <code>GSS_C_GSS_CODE</code> , <i>status_value</i> is a GSS-API status code. If the value is <code>GSS_C_MECH_CODE</code> , then <i>status_value</i> is a mechanism status code.	<i>mech_type</i>	Underlying mechanism that is used to interpret a minor status value. Supply <code>GSS_C_NO_OID</code> to obtain the system default.	<i>message_context</i>	Should be initialized to zero prior to the first call. On return from <code>gss_display_status()</code> , a non-zero <i>status_value</i> parameter indicates that additional messages may be extracted from the status code by means of subsequent calls to <code>gss_display_status()</code> , passing the same <i>status_value</i> , <i>status_type</i> , <i>mech_type</i> , and <i>message_context</i> parameters.	<i>status_string</i>	Textual representation of the <i>status_value</i> . Storage associated with this parameter must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code> .
<i>minor_status</i>	Status code returned by the underlying mechanism.												
<i>status_value</i>	Status value to be converted.												
<i>status_type</i>	If the value is <code>GSS_C_GSS_CODE</code> , <i>status_value</i> is a GSS-API status code. If the value is <code>GSS_C_MECH_CODE</code> , then <i>status_value</i> is a mechanism status code.												
<i>mech_type</i>	Underlying mechanism that is used to interpret a minor status value. Supply <code>GSS_C_NO_OID</code> to obtain the system default.												
<i>message_context</i>	Should be initialized to zero prior to the first call. On return from <code>gss_display_status()</code> , a non-zero <i>status_value</i> parameter indicates that additional messages may be extracted from the status code by means of subsequent calls to <code>gss_display_status()</code> , passing the same <i>status_value</i> , <i>status_type</i> , <i>mech_type</i> , and <i>message_context</i> parameters.												
<i>status_string</i>	Textual representation of the <i>status_value</i> . Storage associated with this parameter must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code> .												
<b>RETURN VALUES</b>	<p>The <code>gss_display_status()</code> function may return the following status codes:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>GSS_S_COMPLETE</code></td> <td>Successful completion.</td> </tr> </table>	<code>GSS_S_COMPLETE</code>	Successful completion.										
<code>GSS_S_COMPLETE</code>	Successful completion.												

## gss\_display\_status(3GSS)

GSS_S_BAD_MECH	Indicates that translation in accordance with an unsupported mechanism type was requested.
GSS_S_BAD_STATUS	The status value was not recognized, or the status type was neither GSS_C_GSS_CODE nor GSS_C_MECH_CODE.
GSS_S_FAILURE	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_acquire_cred(3GSS)`, `gss_release_buffer(3GSS)`, `attributes(5)`

GSS-API Programming Guide

<b>NAME</b>	gss_duplicate_name – create a copy of an internal name								
<b>SYNOPSIS</b>	<pre>cc [flag...] file... -lgss [library...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_duplicate_name</b>(OM_uint32 *minor_status, const     gss_name_t src_name, gss_name_t *dest_name);</pre>								
<b>DESCRIPTION</b>	The <code>gss_duplicate_name()</code> function creates an exact duplicate of the existing internal name <code>src_name</code> . The new <code>dest_name</code> will be independent of the <code>src_name</code> . The <code>src_name</code> and <code>dest_name</code> must both be released, and the release of one does not affect the validity of the other.								
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_duplicate_name()</code> follow:</p> <p><i>minor_status</i>     A mechanism-specific status code.</p> <p><i>src_name</i>         Internal name to be duplicated.</p> <p><i>dest_name</i>         The resultant copy of <code>src_name</code>. Storage associated with this name must be freed by the application after use with a call to <code>gss_release_name(3GSS)</code>.</p>								
<b>RETURN VALUES</b>	<p>The <code>gss_duplicate_name()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>     Successful completion.</p> <p><code>GSS_S_BAD_NAME</code>     The <code>src_name</code> parameter was ill-formed.</p> <p><code>GSS_S_FAILURE</code>     The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <code>minor_status</code> parameter details the error condition.</p>								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWgss (32-bit)</td> </tr> <tr> <td></td> <td>SUNWgssx (64-bit)</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<p><code>gss_release_name(3GSS)</code>, <code>attributes(5)</code></p> <p>GSS-API Programming Guide</p>								

gss\_export\_name(3GSS)

**NAME** | gss\_export\_name – convert a mechanism name to export form

**SYNOPSIS** | 

```
cc [flag ...] file... -lgss [library ...]
#include <gssapi/gssapi.h>

OM_uint32 gss_export_name(OM_uint32 *minor_status, const gss_name_t
input_name, gss_buffer_t exported_name);
```

**DESCRIPTION** | The gss\_export\_name() function allows a GSS-API internal name to be converted into a mechanism-specific name. The function produces a canonical contiguous string representation of a mechanism name, suitable for direct comparison, with memcmp(3C), or for use in authorization functions, matching entries in an access-control list. The input\_name parameter must specify a valid mechanism name, that is, an internal name generated by gss\_accept\_sec\_context(3GSS) or by gss\_canonicalize\_name(3GSS).

**PARAMETERS** | The parameter descriptions for gss\_export\_name() follow:

*minor\_status* | A mechanism-specific status code.

*input\_name* | The mechanism name to be exported.

*exported\_name* | The canonical contiguous string form of *input\_name*. Storage associated with this string must freed by the application after use with gss\_release\_buffer(3GSS).

**RETURN VALUES** | The gss\_export\_name() function may return the following status codes:

GSS\_S\_COMPLETE | Successful completion.

GSS\_S\_NAME\_NOT\_MN | The provided internal name was not a mechanism name.

GSS\_S\_FAILURE | The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the *minor\_status* parameter details the error condition.

**ATTRIBUTES** | See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** | gss\_accept\_sec\_context(3GSS), gss\_canonicalize\_name(3GSS), gss\_release\_buffer(3GSS), memcmp(3C), attributes(5)

*GSS-API Programming Guide*

<b>NAME</b>	gss_export_sec_context – transfer a security context to another process						
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_export_sec_context(OM_uint32 *minor_status,     gss_ctx_id_t *context_handle, gss_buffer_t interprocess_token);</pre>						
<b>DESCRIPTION</b>	<p>The <code>gss_export_sec_context()</code> function generates an interprocess token for transfer to another process within an end system. <code>gss_export_sec_context()</code> and <code>gss_import_sec_context()</code> allow a security context to be transferred between processes on a single machine.</p> <p>The <code>gss_export_sec_context()</code> function supports the sharing of work between multiple processes. This routine is typically used by the context-acceptor, in an application where a single process receives incoming connection requests and accepts security contexts over them, then passes the established context to one or more other processes for message exchange. <code>gss_export_sec_context()</code> deactivates the security context for the calling process and creates an interprocess token which, when passed to <code>gss_import_sec_context()</code> in another process, reactivates the context in the second process. Only a single instantiation of a given context can be active at any one time; a subsequent attempt by a context exporter to access the exported security context will fail.</p> <p>The interprocess token may contain security-sensitive information, for example cryptographic keys. While mechanisms are encouraged to either avoid placing such sensitive information within interprocess tokens or to encrypt the token before returning it to the application, in a typical object-library GSS-API implementation, this might not be possible. Thus, the application must take care to protect the interprocess token and ensure that any process to which the token is transferred is trustworthy. If creation of the interprocess token is successful, the GSS-API deallocates all process-wide resources associated with the security context and sets the <code>context_handle</code> to <code>GSS_C_NO_CONTEXT</code>. In the event of an error that makes it impossible to complete the export of the security context, the function does not return an interprocess token and leaves the security context referenced by the <code>context_handle</code> parameter untouched.</p> <p>Sun's implementation of <code>gss_export_sec_context()</code> does not encrypt the interprocess token. The interprocess token is serialized before it is transferred to another process.</p>						
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_export_sec_context()</code> are as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><i>minor_status</i></td> <td>A mechanism-specific status code.</td> </tr> <tr> <td><i>context_handle</i></td> <td>Context handle identifying the context to transfer.</td> </tr> <tr> <td><i>interprocess_token</i></td> <td>Token to be transferred to target process. Storage associated with this token must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code>.</td> </tr> </table>	<i>minor_status</i>	A mechanism-specific status code.	<i>context_handle</i>	Context handle identifying the context to transfer.	<i>interprocess_token</i>	Token to be transferred to target process. Storage associated with this token must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code> .
<i>minor_status</i>	A mechanism-specific status code.						
<i>context_handle</i>	Context handle identifying the context to transfer.						
<i>interprocess_token</i>	Token to be transferred to target process. Storage associated with this token must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code> .						

`gss_export_sec_context(3GSS)`

**RETURN VALUES** `gss_export_sec_context()` returns one of the following status codes:

<code>GSS_S_COMPLETE</code>	Successful completion.
<code>GSS_S_CONTEXT_EXPIRED</code>	The context has expired.
<code>GSS_S_NO_CONTEXT</code>	The context was invalid.
<code>GSS_S_UNAVAILABLE</code>	The operation is not supported.
<code>GSS_S_FAILURE</code>	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT Level	Safe

**SEE ALSO** `gss_accept_sec_context(3GSS)`, `gss_import_sec_context(3GSS)`, `gss_init_sec_context(3GSS)`, `gss_release_buffer(3GSS)`, `attributes(5)`

*GSS-API Programming Guide*

<b>NAME</b>	gss_get_mic – calculate a cryptographic message
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_get_mic(OM_uint32 *minor_status, const gss_ctx_id_t     context_handle, gss_qop_t qop_req, const gss_buffer_t message_buffer,     gss_buffer_t msg_token);</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_get_mic()</code> function generates a cryptographic MIC for the supplied message, and places the MIC in a token for transfer to the peer application. The <code>qop_req</code> parameter allows a choice between several cryptographic algorithms, if supported by the chosen mechanism.</p> <p>Since some application-level protocols may wish to use tokens emitted by <code>gss_wrap(3GSS)</code> to provide secure framing, the GSS-API allows MICs to be derived from zero-length messages.</p>
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_get_mic()</code> follow:</p> <p><i>minor_status</i>      The status code returned by the underlying mechanism.</p> <p><i>context_handle</i>    Identifies the context on which the message will be sent.</p> <p><i>qop_req</i>            Specifies the requested quality of protection. Callers are encouraged, on portability grounds, to accept the default quality of protection offered by the chosen mechanism, which may be requested by specifying <code>GSS_C_QOP_DEFAULT</code> for this parameter. If an unsupported protection strength is requested, <code>gss_get_mic()</code> will return a <i>major_status</i> of <code>GSS_S_BAD_QOP</code>.</p> <p><i>message_buffer</i>    The message to be protected.</p> <p><i>msg_token</i>         The buffer to receive the token. Storage associated with this message must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code>.</p>
<b>RETURN VALUES</b>	<p><code>gss_get_mic()</code> may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>                      Successful completion.</p> <p><code>GSS_S_CONTEXT_EXPIRED</code>              The context has already expired.</p> <p><code>GSS_S_NO_CONTEXT</code>                    The <i>context_handle</i> parameter did not identify a valid context.</p> <p><code>GSS_S_BAD_QOP</code>                        The specified QOP is not supported by the mechanism.</p> <p><code>GSS_S_FAILURE</code>                        The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>

gss\_get\_mic(3GSS)

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_release_buffer(3GSS)`, `gss_wrap(3GSS)`, `attributes(5)`  
GSS-API Programming Guide



<b>NAME</b>	gss_import_name – convert a contiguous string name to GSS_API internal format
<b>SYNOPSIS</b>	<pre>cc [flag...] file... -lgss [library...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_import_name(OM_uint32 * minor_status, const     gss_buffer_t input_name_buffer, const gss_OID input_name_type,     gss_name_t *output_name);</pre>
<b>DESCRIPTION</b>	The <code>gss_import_name()</code> function converts a contiguous string name to internal form. In general, the internal name returned by means of the <code>output_name</code> parameter will not be a mechanism name; the exception to this is if the <code>input_name_type</code> indicates that the contiguous string provided by means of the <code>input_name_buffer</code> parameter is of type <code>GSS_C_NT_EXPORT_NAME</code> , in which case, the returned internal name will be a mechanism name for the mechanism that exported the name.
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_import_name()</code> follow:</p> <p><i>minor_status</i>            Status code returned by the underlying mechanism.</p> <p><i>input_name_buffer</i>        The <code>gss_buffer_desc</code> structure containing the name to be imported. The application must allocate this explicitly. This argument must be deallocated with <code>gss_release_buffer(3GSS)</code> when the application is done with it.</p> <p><i>input_name_type</i>          A <code>gss_OID</code> that specifies the format that the <code>input_name_buffer</code> is in.</p> <p><i>output_name</i>              The <code>gss_name_t</code> structure to receive the name.</p>
<b>RETURN VALUES</b>	<p>The <code>gss_import_name()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>            The <code>gss_import_name()</code> function completed successfully.</p> <p><code>GSS_S_BAD_NAME_TYPE</code>        The <code>input_name_type</code> was unrecognized.</p> <p><code>GSS_S_BAD_NAME</code>            The <code>input_name</code> parameter could not be interpreted as a name of the specified type.</p> <p><code>GSS_S_BAD_MECH</code>            The <code>input_name_type</code> was <code>GSS_C_NT_EXPORT_NAME</code>, but the mechanism contained within the <code>input_name</code> is not supported.</p> <p><code>GSS_S_FAILURE</code>            The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <code>minor_status</code> parameter details the error condition.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

gss\_import\_name(3GSS)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** [gss\\_release\\_buffer\(3GSS\)](#), [attributes\(5\)](#)

[GSS-API Programming Guide](#)

gss\_import\_sec\_context(3GSS)

**NAME** | gss\_import\_sec\_context – import security context established by another process

**SYNOPSIS** | 

```
cc -flag ... file...-lgss [library ...]
#include <gssapi/gssapi.h>

OM_uint32 gss_import_sec_context(OM_uint32 *minor_status, const
    gss_buffer_t interprocess_token, gss_ctx_id_t *context_handle);
```

**DESCRIPTION** | The `gss_import_sec_context()` function allows a process to import a security context established by another process. A given interprocess token can be imported only once. See `gss_export_sec_context(3GSS)`.

**PARAMETERS** | The parameter descriptions for `gss_import_sec_context()` are as follows:

<i>minor_status</i>	A mechanism-specific status code.
<i>interprocess_token</i>	Token received from exporting process.
<i>context_handle</i>	Context handle of newly reactivated context. Resources associated with this context handle must be released by the application after use with a call to <code>gss_delete_sec_context(3GSS)</code> .

**RETURN VALUES** | `gss_import_sec_context()` returns one of the following status codes:

GSS_S_COMPLETE	Successful completion.
GSS_S_NO_CONTEXT	The token did not contain a valid context reference.
GSS_S_DEFECTIVE_TOKEN	The token was invalid.
GSS_S_UNAVAILABLE	The operation is unavailable.
GSS_S_UNAUTHORIZED	Local policy prevents the import of this context by the current process.
GSS_S_FAILURE	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.

**ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT Level	Safe

`gss_import_sec_context(3GSS)`

**SEE ALSO** | `gss_accept_sec_context(3GSS)`, `gss_context_time(3GSS)`,  
`gss_delete_sec_context(3GSS)`, `gss_export_sec_content(3GSS)`,  
`gss_init_sec_context(3GSS)`, `gss_inquire_context(3GSS)`,  
`gss_process_context_token(3GSS)`, `gss_wrap_size_limit(3GSS)`,  
`attributes(5)`

*GSS-API Programming Guide*

<b>NAME</b>	gss_indicate_mechs – determine available security mechanisms								
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_indicate_mechs(OM_uint32 *minor_status, gss_OID_set     *mech_set);</pre>								
<b>DESCRIPTION</b>	The <code>gss_indicate_mechs()</code> function enables an application to determine available underlying security mechanisms.								
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_indicate_mechs()</code> follow:</p> <p><i>minor_status</i>     A mechanism-specific status code.</p> <p><i>mech_set</i>         Set of supported mechanisms. The returned <code>gss_OID_set</code> value will be a dynamically-allocated OID set that should be released by the caller after use with a call to <code>gss_release_oid_set(3GSS)</code>.</p>								
<b>RETURN VALUES</b>	<p>The <code>gss_indicate_mechs()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>     Successful completion.</p> <p><code>GSS_S_FAILURE</code>     The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWgss (32-bit)</td> </tr> <tr> <td></td> <td>SUNWgssx (64-bit)</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<p><code>gss_release_oid_set(3GSS)</code>, <code>attributes(5)</code></p> <p>GSS-API Programming Guide</p>								

## gss\_init\_sec\_context(3GSS)

<b>NAME</b>	<code>gss_init_sec_context</code> – initiate a GSS-API security context with a peer application
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_init_sec_context</b>(OM_uint32 *minor_status, const     gss_cred_id_t initiator_cred_handle, gss_ctx_id_t *context_handle,     const gss_name_t *target_name, const gss_OID mech_type, OM_uint32     req_flags, OM_uint32 time_req, const gss_channel_bindings_t     input_chan_bindings, const gss_buffer_t input_token, gss_OID     *actual_mech_type, gss_buffer_t output_token, OM_uint32 *ret_flags,     OM_uint32 *time_rec);</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_init_sec_context()</code> function initiates the establishment of a security context between the application and a remote peer. Initially, the <code>input_token</code> parameter should be specified either as <code>GSS_C_NO_BUFFER</code>, or as a pointer to a <code>gss_buffer_desc</code> object with a length field that contains a zero value. The routine may return a <code>output_token</code>, which should be transferred to the peer application, which will present it to <code>gss_accept_sec_context(3GSS)</code>. If no token need be sent, <code>gss_init_sec_context()</code> will indicate this by setting the length field of the <code>output_token</code> argument to zero. To complete context establishment, one or more reply tokens may be required from the peer application; if so, <code>gss_init_sec_context()</code> will return a status code that contains the supplementary information bit <code>GSS_S_CONTINUE_NEEDED</code>. In this case, make another call to <code>gss_init_sec_context()</code> when the reply token is received from the peer application and pass the reply token to <code>gss_init_sec_context()</code> by means of the <code>input_token</code> parameter.</p> <p>Construct portable applications to use the token length and return status to determine whether to send or wait for a token.</p> <p>Whenever the routine returns a major status that includes the value <code>GSS_S_CONTINUE_NEEDED</code>, the context is not fully established, and the following restrictions apply to the output parameters:</p> <ul style="list-style-type: none"><li>■ The value returned by means of the <code>time_rec</code> parameter is undefined. Unless the accompanying <code>ret_flags</code> parameter contains the bit <code>GSS_C_PROT_READY_FLAG</code>, which indicates that per-message services may be applied in advance of a successful completion status, the value returned by means of the <code>actual_mech_type</code> parameter is undefined until the routine returns a major status value of <code>GSS_S_COMPLETE</code>.</li><li>■ The values of the <code>GSS_C_DELEG_FLAG</code>, <code>GSS_C_MUTUAL_FLAG</code>, <code>GSS_C_REPLAY_FLAG</code>, <code>GSS_C_SEQUENCE_FLAG</code>, <code>GSS_C_CONF_FLAG</code>, <code>GSS_C_INTEG_FLAG</code> and <code>GSS_C_ANON_FLAG</code> bits returned by the <code>ret_flags</code> parameter contain values that will be valid if context establishment succeeds. For example, if the application requests a service such as delegation or anonymous authentication by means of the <code>req_flags</code> argument, and the service is unavailable from the underlying mechanism, <code>gss_init_sec_context()</code> generates a token that will not provide the service, and it indicate by means of the <code>ret_flags</code> argument that the service will not be supported. The application may choose to abort context</li></ul>

`gss_init_sec_context(3GSS)`

establishment by calling `gss_delete_sec_context(3GSS)` if it cannot continue without the service, or if the service was merely desired but not mandatory, it may transmit the token and continue context establishment.

- The values of the `GSS_C_PROT_READY_FLAG` and `GSS_C_TRANS_FLAG` bits within `ret_flags` indicate the actual state at the time `gss_init_sec_context()` returns, whether or not the context is fully established.
- The GSS-API sets the `GSS_C_PROT_READY_FLAG` in the final `ret_flags` returned to a caller, for example, when accompanied by a `GSS_S_COMPLETE` status code. However, applications should not rely on this behavior, as the flag was not defined in Version 1 of the GSS-API. Instead, applications should determine what per-message services are available after a successful context establishment according to the `GSS_C_INTEG_FLAG` and `GSS_C_CONF_FLAG` values.
- All other bits within the `ret_flags` argument are set to zero.

If the initial call of `gss_init_sec_context()` fails, the GSS-API does not create a context object; it leaves the value of the `context_handle` parameter set to `GSS_C_NO_CONTEXT` to indicate this. In the event of failure on a subsequent call, the GSS-API leaves the security context untouched for the application to delete using `gss_delete_sec_context(3GSS)`.

During context establishment, the informational status bits `GSS_S_OLD_TOKEN` and `GSS_S_DUPLICATE_TOKEN` indicate fatal errors, and GSS-API mechanisms should always return them in association with a status code of `GSS_S_FAILURE`. This pairing requirement was not part of Version 1 of the GSS-API specification, so applications that wish to run on Version 1 implementations must special-case these codes.

## PARAMETERS

The parameter descriptions for `gss_init_sec_context()` follow:

<i>minor_status</i>	A mechanism specific status code.
<i>initiator_cred_handle</i>	The handle for the credentials claimed. Supply <code>GSS_C_NO_CREDENTIAL</code> to act as a default initiator principal. If no default initiator is defined, the function returns <code>GSS_S_NO_CRED</code> .
<i>context_handle</i>	The context handle for a new context. Supply the value <code>GSS_C_NO_CONTEXT</code> for the first call, and use the value returned in any continuation calls. The resources associated with <code>context_handle</code> must be released by the application after use by a call to <code>gss_delete_sec_context(3GSS)</code> .
<i>target_name</i>	The name of the target.
<i>mech_type</i>	The object ID of the desired mechanism. To obtain a specific default, supply the value <code>GSS_C_NO_ID</code> .
<i>req_flags</i>	Contains independent flags, each of which will request that the context support a specific service option. A symbolic name is provided for each flag. Logically-OR the symbolic name to

## gss\_init\_sec\_context(3GSS)

	<p>the corresponding required flag to form the bit-mask value. <i>req_flags</i> may contain one of the following values:</p> <p><b>GSS_C_DELEG_FLAG</b> If true, delegate credentials to a remote peer. Do not delegate the credentials if the value is false.</p> <p><b>GSS_C_MUTUAL_FLAG</b> If true, request that the peer authenticate itself. If false, authenticate to the remote peer only.</p> <p><b>GSS_C_REPLAY_FLAG</b> If true, enable replay detection for messages protected with <code>gss_wrap(3GSS)</code> or <code>gss_get_mic(3GSS)</code>. Do not attempt to detect replayed messages if false.</p> <p><b>GSS_C_SEQUENCE_FLAG</b> If true, enable detection of out-of-sequence protected messages. Do not attempt to detect out-of-sequence messages if false.</p> <p><b>GSS_C_CONF_FLAG</b> If true, request that confidential service be made available by means of <code>gss_wrap(3GSS)</code>. If false, no per-message confidential service is required.</p> <p><b>GSS_C_INTEG_FLAG</b> If true, request that integrity service be made available by means of <code>gss_wrap(3GSS)</code> or <code>gss_get_mic(3GSS)</code>. If false, no per-message integrity service is required.</p> <p><b>GSS_C_ANON_FLAG</b> If true, do not reveal the initiator's identify to the acceptor. If false, authenticate normally.</p>
<i>time_req</i>	The number of seconds for which the context will remain valid. Supply a zero value to <i>time_req</i> to request a default validity period.
<i>input_chan_bindings</i>	Optional application-specified bindings. Allows application to securely bind channel identification information to the security context. Set to <code>GSS_C_NO_CHANNEL_BINDINGS</code> if you do not want to use channel bindings.
<i>input_token</i>	Token received from the peer application. On the initial call, supply <code>GSS_C_NO_BUFFER</code> or a pointer to a buffer containing the value <code>GSS_C_EMPTY_BUFFER</code> .
<i>actual_mech_type</i>	The actual mechanism used. The OID returned by means of this parameter will be pointer to static storage that should be treated as read-only. The application should not attempt to free it. To obtain a specific default, supply the value <code>GSS_C_NO_ID</code> . Specify <code>NULL</code> if the parameter is not required.



gss\_init\_sec\_context(3GSS)

*output\_token*

The token to send to the peer application. If the length field of the returned buffer is zero, no token need be sent to the peer application. After use storage associated with this buffer must be freed by the application by a call to `gss_release_buffer(3GSS)`.

*ret\_flags*

Contains various independent flags, each of which indicates that the context supports a specific service option. If not needed, specify `NULL`. Test the returned bit-mask *ret\_flags* value against its symbolic name to determine if the given option is supported by the context. *ret\_flags* may contain one of the following values:

`GSS_C_DELEG_FLAG`

If true, credentials were delegated to the remote peer. If false, no credentials were delegated.

`GSS_C_MUTUAL_FLAG`

If true, the remote peer authenticated itself. If false, the remote peer did not authenticate itself.

`GSS_C_REPLY_FLAG`

If true, replay of protected messages will be detected. If false, replayed messages will not be detected.

`GSS_C_SEQUENCE_FLAG`

If true, out of sequence protected messages will be detected. If false, they will not be detected.

`GSS_C_CONF_FLAG`

If true, confidential service may be invoked by calling the `gss_wrap()` routine. If false, no confidentiality service is available by means of `gss_wrap(3GSS)`. `gss_wrap()` will provide message encapsulation, data-origin authentication and integrity services only.

`GSS_C_INTEG_FLAG`

If true, integrity service may be invoked by calling either the `gss_wrap(3GSS)` or `gss_get_mic(3GSS)` routine. If false, per-message integrity service is not available.

`GSS_C_ANON_FLAG`

If true, the initiator's identity has not been revealed; it will not be revealed if any emitted token is passed to the acceptor. If false, the initiator has been or will be authenticated normally.

`GSS_C_PROT_READY_FLAG`

If true, the protection services specified by the states of `GSS_C_CONF_FLAG` and `GSS_C_INTEG_FLAG` are available if the accompanying major status return value is either `GSS_S_COMPLETE` or `GSS_S_CONTINUE_NEEDED`. If false,

## gss\_init\_sec\_context(3GSS)

the protection services are available only if the accompanying major status return value is `GSS_S_COMPLETE`.

### `GSS_C_TRANS_FLAG`

If true, the resultant security context may be transferred to other processes by means of a call to `gss_export_sec_context(3GSS)`. If false, the security context cannot be transferred.

### `time_rec`

The number of seconds for which the context will remain valid. Specify `NULL` if the parameter is not required.

## RETURN VALUES

`gss_init_sec_context()` may return the following status codes:

<code>GSS_S_COMPLETE</code>	Successful completion.
<code>GSS_S_CONTINUE_NEEDED</code>	A token from the peer application is required to complete the context, and <code>gss_init_sec_context()</code> must be called again with that token.
<code>GSS_S_DEFECTIVE_TOKEN</code>	Consistency checks performed on the <i>input_token</i> failed.
<code>GSS_S_DEFECTIVE_CREDENTIAL</code>	Consistency checks performed on the credential failed.
<code>GSS_S_NO_CRED</code>	The supplied credentials are not valid for context acceptance, or the credential handle does not reference any credentials.
<code>GSS_S_CREDENTIALS_EXPIRED</code>	The referenced credentials have expired.
<code>GSS_S_BAD_BINDINGS</code>	The <i>input_token</i> contains different channel bindings than those specified by means of the <i>input_chan_bindings</i> parameter.
<code>GSS_S_BAD_SIG</code>	The <i>input_token</i> contains an invalid MIC or a MIC that cannot be verified.
<code>GSS_S_OLD_TOKEN</code>	The <i>input_token</i> is too old. This is a fatal error while establishing context.
<code>GSS_S_DUPLICATE_TOKEN</code>	The <i>input_token</i> is valid, but it is a duplicate of a token already processed. This is a fatal error while establishing context.
<code>GSS_S_NO_CONTEXT</code>	The supplied context handle does not refer to a valid context.
<code>GSS_S_BAD_NAME_TYPE</code>	The provided <i>target_name</i> parameter contains an invalid or unsupported <i>name</i> type.

`gss_init_sec_context(3GSS)`

<code>GSS_S_BAD_NAME</code>	The supplied <i>target_name</i> parameter is ill-formed.
<code>GSS_S_BAD_MECH</code>	The token received specifies a mechanism that is not supported by the implementation or the provided credential.
<code>GSS_S_FAILURE</code>	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.

**EXAMPLES** **EXAMPLE 1** Invoking `gss_init_sec_context()` Within a Loop

A typical portable caller should always invoke `gss_init_sec_context()` within a loop:

```
int context_established = 0;
gss_ctx_id_t context_hdl = GSS_C_NO_CONTEXT;
...
input_token->length = 0;

while (!context_established) {
    maj_stat = gss_init_sec_context(&min_stat,
                                   cred_hdl,
                                   &context_hdl,
                                   target_name,
                                   desired_mech,
                                   desired_services,
                                   desired_time,
                                   input_bindings,
                                   input_token,
                                   &actual_mech,
                                   output_token,
                                   &actual_services,
                                   &actual_time);

    if (GSS_ERROR(maj_stat)) {
        report_error(maj_stat, min_stat);
    };

    if (output_token->length != 0) {
        send_token_to_peer(output_token);
        gss_release_buffer(&min_stat, output_token)
    };
    if (GSS_ERROR(maj_stat)) {

        if (context_hdl != GSS_C_NO_CONTEXT)
            gss_delete_sec_context(&min_stat,
                                   &context_hdl,
                                   GSS_C_NO_BUFFER);

        break;
    };
    if (maj_stat & GSS_S_CONTINUE_NEEDED) {
        receive_token_from_peer(input_token);
    }
}
```

## gss\_init\_sec\_context(3GSS)

**EXAMPLE 1** Invoking `gss_init_sec_context()` Within a Loop *(Continued)*

```
    } else {  
        context_established = 1;  
    };  
};
```

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_delete_sec_context(3GSS)`, `gss_export_sec_context(3GSS)`,  
`gss_get_mic(3GSS)`, `gss_wrap(3GSS)`, `attributes(5)`

GSS-API Programming Guide

<b>NAME</b>	gss_inquire_context – obtain information about a security context														
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_inquire_context</b>(OM_uint32 *minor_status, const     gss_ctx_id_t context_handle, gss_name_t *src_name, gss_name_t     *targ_name, OM_uint32 *lifetime_rec, gss_OID *mech_type, OM_uint32     *ctx_flags, int *locally_initiated, int *open);</pre>														
<b>DESCRIPTION</b>	The <code>gss_inquire_context()</code> function obtains information about a security context. The caller must already have obtained a handle that refers to the context, although the context need not be fully established.														
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_inquire_context()</code> are as follows:</p> <table border="0"> <tr> <td style="vertical-align: top;"><i>minor_status</i></td> <td>A mechanism-specific status code.</td> </tr> <tr> <td style="vertical-align: top;"><i>context_handle</i></td> <td>A handle that refers to the security context.</td> </tr> <tr> <td style="vertical-align: top;"><i>src_name</i></td> <td>The name of the context initiator. If the context was established using anonymous authentication, and if the application invoking <code>gss_inquire_context()</code> is the context acceptor, an anonymous name is returned. Storage associated with this name must be freed by the application after use with a call to <code>gss_release_name()</code>. Specify NULL if the parameter is not required.</td> </tr> <tr> <td style="vertical-align: top;"><i>targ_name</i></td> <td>The name of the context acceptor. Storage associated with this name must be freed by the application after use with a call to <code>gss_release_name()</code>. If the context acceptor did not authenticate itself, and if the initiator did not specify a target name in its call to <code>gss_init_sec_context()</code>, the value <code>GSS_C_NO_NAME</code> is returned. Specify NULL if the parameter is not required.</td> </tr> <tr> <td style="vertical-align: top;"><i>lifetime_rec</i></td> <td>The number of seconds for which the context will remain valid. If the context has expired, this parameter will be set to zero. Specify NULL if the parameter is not required.</td> </tr> <tr> <td style="vertical-align: top;"><i>mech_type</i></td> <td>The security mechanism providing the context. The returned OID is a pointer to static storage that should be treated as read-only by the application; in particular, the application should not attempt to free it. Specify NULL if the parameter is not required.</td> </tr> <tr> <td style="vertical-align: top;"><i>ctx_flags</i></td> <td>Contains various independent flags, each of which indicates that the context supports (or is expected to support, if <code>ctx_open</code> is false) a specific service option. If not needed, specify NULL. Symbolic names are provided for each flag, and the symbolic names corresponding to the required flags should be logically ANDed with the <code>ret_flags</code> value to test whether a given option is supported by the context. The flags are:</td> </tr> </table>	<i>minor_status</i>	A mechanism-specific status code.	<i>context_handle</i>	A handle that refers to the security context.	<i>src_name</i>	The name of the context initiator. If the context was established using anonymous authentication, and if the application invoking <code>gss_inquire_context()</code> is the context acceptor, an anonymous name is returned. Storage associated with this name must be freed by the application after use with a call to <code>gss_release_name()</code> . Specify NULL if the parameter is not required.	<i>targ_name</i>	The name of the context acceptor. Storage associated with this name must be freed by the application after use with a call to <code>gss_release_name()</code> . If the context acceptor did not authenticate itself, and if the initiator did not specify a target name in its call to <code>gss_init_sec_context()</code> , the value <code>GSS_C_NO_NAME</code> is returned. Specify NULL if the parameter is not required.	<i>lifetime_rec</i>	The number of seconds for which the context will remain valid. If the context has expired, this parameter will be set to zero. Specify NULL if the parameter is not required.	<i>mech_type</i>	The security mechanism providing the context. The returned OID is a pointer to static storage that should be treated as read-only by the application; in particular, the application should not attempt to free it. Specify NULL if the parameter is not required.	<i>ctx_flags</i>	Contains various independent flags, each of which indicates that the context supports (or is expected to support, if <code>ctx_open</code> is false) a specific service option. If not needed, specify NULL. Symbolic names are provided for each flag, and the symbolic names corresponding to the required flags should be logically ANDed with the <code>ret_flags</code> value to test whether a given option is supported by the context. The flags are:
<i>minor_status</i>	A mechanism-specific status code.														
<i>context_handle</i>	A handle that refers to the security context.														
<i>src_name</i>	The name of the context initiator. If the context was established using anonymous authentication, and if the application invoking <code>gss_inquire_context()</code> is the context acceptor, an anonymous name is returned. Storage associated with this name must be freed by the application after use with a call to <code>gss_release_name()</code> . Specify NULL if the parameter is not required.														
<i>targ_name</i>	The name of the context acceptor. Storage associated with this name must be freed by the application after use with a call to <code>gss_release_name()</code> . If the context acceptor did not authenticate itself, and if the initiator did not specify a target name in its call to <code>gss_init_sec_context()</code> , the value <code>GSS_C_NO_NAME</code> is returned. Specify NULL if the parameter is not required.														
<i>lifetime_rec</i>	The number of seconds for which the context will remain valid. If the context has expired, this parameter will be set to zero. Specify NULL if the parameter is not required.														
<i>mech_type</i>	The security mechanism providing the context. The returned OID is a pointer to static storage that should be treated as read-only by the application; in particular, the application should not attempt to free it. Specify NULL if the parameter is not required.														
<i>ctx_flags</i>	Contains various independent flags, each of which indicates that the context supports (or is expected to support, if <code>ctx_open</code> is false) a specific service option. If not needed, specify NULL. Symbolic names are provided for each flag, and the symbolic names corresponding to the required flags should be logically ANDed with the <code>ret_flags</code> value to test whether a given option is supported by the context. The flags are:														

## gss\_inquire\_context(3GSS)

### GSS\_C\_DELEG\_FLAG

If true, credentials were delegated from the initiator to the acceptor. If false, no credentials were delegated.

### GSS\_C\_MUTUAL\_FLAG

If true, the acceptor was authenticated to the initiator. If false, the acceptor did not authenticate itself.

### GSS\_C\_REPLAY\_FLAG

If true, the replay of protected messages will be detected. If false, replayed messages will not be detected.

### GSS\_C\_SEQUENCE\_FLAG

If true, out-of-sequence protected messages will be detected. If false, out-of-sequence messages will not be detected.

### GSS\_C\_CONF\_FLAG

If true, confidential service may be invoked by calling the `gss_wrap(3GSS)` routine. If false, no confidential service is available through `gss_wrap()`. `gss_wrap()` provides message encapsulation, data-origin authentication, and integrity services only.

### GSS\_C\_INTEG\_FLAG

If true, integrity service can be invoked by calling either the `gss_get_mic()` or the `gss_wrap()` routine. If false, per-message integrity service is unavailable.

### GSS\_C\_ANON\_FLAG

If true, the initiator's identity is not revealed to the acceptor. The `src_name` parameter, if requested, contains an anonymous internal name. If false, the initiator has been authenticated normally.

### GSS\_C\_PROT\_READY\_FLAG

If true, the protection services, as specified by the states of the `GSS_C_CONF_FLAG` and `GSS_C_INTEG_FLAG`, are available for use. If false, they are available only if the context is fully established, that is, if the `open` parameter is non-zero.

### GSS\_C\_TRANS\_FLAG

If true, resultant security context can be transferred to other processes through a call to `gss_export_sec_context()`. If false, the security context is not transferable.

*locally\_initiated*

Non-zero if the invoking application is the context initiator. Specify NULL if the parameter is not required.

*open*

Non-zero if the context is fully established; zero if a context-establishment token is expected from the peer application. Specify NULL if the parameter is not required.

**RETURN VALUES** `gss_inquire_context()` returns one of the following status codes:

`gss_inquire_context(3GSS)`

`GSS_S_COMPLETE` Successful completion.

`GSS_S_NO_CONTEXT` The referenced context could not be accessed.

`GSS_S_FAILURE` The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the *minor\_status* parameter details the error condition.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_accept_sec_context(3GSS)`, `gss_context_time(3GSS)`,  
`gss_delete_sec_context(3GSS)`, `gss_export_sec_context(3GSS)`,  
`gss_import_sec_context(3GSS)`, `gss_init_sec_context(3GSS)`,  
`gss_process_context_token(3GSS)`, `gss_wrap(3GSS)`,  
`gss_wrap_size_limit(3GSS)`, `attributes(5)`

*GSS-API Programming Guide*

## gss\_inquire\_cred(3GSS)

<b>NAME</b>	<code>gss_inquire_cred</code> – obtain information about a credential												
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_inquire_cred</b>(OM_uint32 *minor_status, const     gss_cred_id_t cred_handle, gss_name_t *name, OM_uint32 *lifetime,     gss_cred_usage_t *cred_usage, gss_OID_set *mechanisms);</pre>												
<b>DESCRIPTION</b>	Use the <code>gss_inquire_cred()</code> function to obtain information about a credential.												
<b>PARAMETERS</b>	The parameter descriptions for <code>gss_acquire_cred()</code> follow:  <table><tr><td><i>minor_status</i></td><td>A mechanism specific status code.</td></tr><tr><td><i>cred_handle</i></td><td>A handle that refers to the target credential. Specify <code>GSS_C_NO_CREDENTIAL</code> to inquire about the default initiator principal.</td></tr><tr><td><i>name</i></td><td>The name whose identity the credential asserts. Any storage associated with this name should be freed by the application after use by a call to <code>gss_release_name(3GSS)</code>.</td></tr><tr><td><i>lifetime</i></td><td>The number of seconds for which the credential will remain valid. If the credential has expired, this parameter will be set to zero. Specify <code>NULL</code> if this parameter is not required.</td></tr><tr><td><i>cred_usage</i></td><td>How the credential may be used. The <i>cred_usage</i> parameter may contain one of the following values: <code>GSS_C_INITIATE</code>, <code>GSS_C_ACCEPT</code>, or <code>GSS_C_BOTH</code>. Specify <code>NULL</code> if this parameter is not required.</td></tr><tr><td><i>mechanisms</i></td><td>The set of mechanisms which the credential supports. Storage for the returned OID-set must be freed by the application after use by a call to <code>gss_release_oid_set(3GSS)</code>. Specify <code>NULL</code> if this parameter is not required.</td></tr></table>	<i>minor_status</i>	A mechanism specific status code.	<i>cred_handle</i>	A handle that refers to the target credential. Specify <code>GSS_C_NO_CREDENTIAL</code> to inquire about the default initiator principal.	<i>name</i>	The name whose identity the credential asserts. Any storage associated with this name should be freed by the application after use by a call to <code>gss_release_name(3GSS)</code> .	<i>lifetime</i>	The number of seconds for which the credential will remain valid. If the credential has expired, this parameter will be set to zero. Specify <code>NULL</code> if this parameter is not required.	<i>cred_usage</i>	How the credential may be used. The <i>cred_usage</i> parameter may contain one of the following values: <code>GSS_C_INITIATE</code> , <code>GSS_C_ACCEPT</code> , or <code>GSS_C_BOTH</code> . Specify <code>NULL</code> if this parameter is not required.	<i>mechanisms</i>	The set of mechanisms which the credential supports. Storage for the returned OID-set must be freed by the application after use by a call to <code>gss_release_oid_set(3GSS)</code> . Specify <code>NULL</code> if this parameter is not required.
<i>minor_status</i>	A mechanism specific status code.												
<i>cred_handle</i>	A handle that refers to the target credential. Specify <code>GSS_C_NO_CREDENTIAL</code> to inquire about the default initiator principal.												
<i>name</i>	The name whose identity the credential asserts. Any storage associated with this name should be freed by the application after use by a call to <code>gss_release_name(3GSS)</code> .												
<i>lifetime</i>	The number of seconds for which the credential will remain valid. If the credential has expired, this parameter will be set to zero. Specify <code>NULL</code> if this parameter is not required.												
<i>cred_usage</i>	How the credential may be used. The <i>cred_usage</i> parameter may contain one of the following values: <code>GSS_C_INITIATE</code> , <code>GSS_C_ACCEPT</code> , or <code>GSS_C_BOTH</code> . Specify <code>NULL</code> if this parameter is not required.												
<i>mechanisms</i>	The set of mechanisms which the credential supports. Storage for the returned OID-set must be freed by the application after use by a call to <code>gss_release_oid_set(3GSS)</code> . Specify <code>NULL</code> if this parameter is not required.												
<b>RETURN VALUES</b>	<code>gss_acquire_cred()</code> may return the following status codes:  <table><tr><td><code>GSS_S_COMPLETE</code></td><td>Successful completion.</td></tr><tr><td><code>GSS_S_NO_CRED</code></td><td>The referenced credentials could not be accessed.</td></tr><tr><td><code>GSS_S_DEFECTIVE_CREDENTIAL</code></td><td>The referenced credentials were invalid.</td></tr><tr><td><code>GSS_S_CREDENTIALS_EXPIRED</code></td><td>The referenced credentials have expired. If the <i>lifetime</i> parameter was not passed as <code>NULL</code>, it will be set to 0.</td></tr><tr><td><code>GSS_S_FAILURE</code></td><td>The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i></td></tr></table>	<code>GSS_S_COMPLETE</code>	Successful completion.	<code>GSS_S_NO_CRED</code>	The referenced credentials could not be accessed.	<code>GSS_S_DEFECTIVE_CREDENTIAL</code>	The referenced credentials were invalid.	<code>GSS_S_CREDENTIALS_EXPIRED</code>	The referenced credentials have expired. If the <i>lifetime</i> parameter was not passed as <code>NULL</code> , it will be set to 0.	<code>GSS_S_FAILURE</code>	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i>		
<code>GSS_S_COMPLETE</code>	Successful completion.												
<code>GSS_S_NO_CRED</code>	The referenced credentials could not be accessed.												
<code>GSS_S_DEFECTIVE_CREDENTIAL</code>	The referenced credentials were invalid.												
<code>GSS_S_CREDENTIALS_EXPIRED</code>	The referenced credentials have expired. If the <i>lifetime</i> parameter was not passed as <code>NULL</code> , it will be set to 0.												
<code>GSS_S_FAILURE</code>	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i>												



`gss_inquire_cred(3GSS)`

parameter details the error condition.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_release_name(3GSS)`, `gss_release_oid_set(3GSS)`, `attributes(5)`  
GSS-API Programming Guide

## gss\_inquire\_cred\_by\_mech(3GSS)

<b>NAME</b>	gss_inquire_cred_by_mech – obtain per-mechanism information about a credential														
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_inquire_cred_by_mech</b>(OM_uint32 *minor_status, const     gss_cred_id_t cred_handle, const gss_OID mech_type, gss_name_t     *name, OM_uint32 *initiator_lifetime, OM_uint32 *acceptor_lifetime,     gss_cred_usage_t *cred_usage);</pre>														
<b>DESCRIPTION</b>	The <code>gss_inquire_cred_by_mech()</code> obtains per-mechanism information about a credential.														
<b>PARAMETERS</b>	The parameter descriptions for <code>gss_inquire_cred_by_mech()</code> follow: <table><tr><td><i>minor_status</i></td><td>A mechanism specific status code.</td></tr><tr><td><i>cred_handle</i></td><td>A handle that refers to the target credential. Specify <code>GSS_C_NO_CREDENTIAL</code> to inquire about the default initiator principal.</td></tr><tr><td><i>mech_type</i></td><td>The mechanism for which the information should be returned.</td></tr><tr><td><i>name</i></td><td>The name whose identity the credential asserts. Any storage associated with this <i>name</i> must be freed by the application after use by a call to <code>gss_release_name(3GSS)</code>.</td></tr><tr><td><i>initiator_lifetime</i></td><td>The number of seconds that the credential is capable of initiating security contexts under the specified mechanism. If the credential can no longer be used to initiate contexts, or if the credential usage for this mechanism is <code>GSS_C_ACCEPT</code>, this parameter will be set to 0. Specify <code>NULL</code> if this parameter is not required.</td></tr><tr><td><i>acceptor_lifetime</i></td><td>The number of seconds that the credential is capable of accepting security contexts under the specified mechanism. If the credential can no longer be used to accept contexts, or if the credential usage for this mechanism is <code>GSS_C_INITIATE</code>, this parameter will be set to 0. Specify <code>NULL</code> if this parameter is not required.</td></tr><tr><td><i>cred_usage</i></td><td>How the credential may be used with the specified mechanism. The <i>cred_usage</i> parameter may contain one of the following values: <code>GSS_C_INITIATE</code>, <code>GSS_C_ACCEPT</code>, or <code>GSS_C_BOTH</code>. Specify <code>NULL</code> if this parameter is not required.</td></tr></table>	<i>minor_status</i>	A mechanism specific status code.	<i>cred_handle</i>	A handle that refers to the target credential. Specify <code>GSS_C_NO_CREDENTIAL</code> to inquire about the default initiator principal.	<i>mech_type</i>	The mechanism for which the information should be returned.	<i>name</i>	The name whose identity the credential asserts. Any storage associated with this <i>name</i> must be freed by the application after use by a call to <code>gss_release_name(3GSS)</code> .	<i>initiator_lifetime</i>	The number of seconds that the credential is capable of initiating security contexts under the specified mechanism. If the credential can no longer be used to initiate contexts, or if the credential usage for this mechanism is <code>GSS_C_ACCEPT</code> , this parameter will be set to 0. Specify <code>NULL</code> if this parameter is not required.	<i>acceptor_lifetime</i>	The number of seconds that the credential is capable of accepting security contexts under the specified mechanism. If the credential can no longer be used to accept contexts, or if the credential usage for this mechanism is <code>GSS_C_INITIATE</code> , this parameter will be set to 0. Specify <code>NULL</code> if this parameter is not required.	<i>cred_usage</i>	How the credential may be used with the specified mechanism. The <i>cred_usage</i> parameter may contain one of the following values: <code>GSS_C_INITIATE</code> , <code>GSS_C_ACCEPT</code> , or <code>GSS_C_BOTH</code> . Specify <code>NULL</code> if this parameter is not required.
<i>minor_status</i>	A mechanism specific status code.														
<i>cred_handle</i>	A handle that refers to the target credential. Specify <code>GSS_C_NO_CREDENTIAL</code> to inquire about the default initiator principal.														
<i>mech_type</i>	The mechanism for which the information should be returned.														
<i>name</i>	The name whose identity the credential asserts. Any storage associated with this <i>name</i> must be freed by the application after use by a call to <code>gss_release_name(3GSS)</code> .														
<i>initiator_lifetime</i>	The number of seconds that the credential is capable of initiating security contexts under the specified mechanism. If the credential can no longer be used to initiate contexts, or if the credential usage for this mechanism is <code>GSS_C_ACCEPT</code> , this parameter will be set to 0. Specify <code>NULL</code> if this parameter is not required.														
<i>acceptor_lifetime</i>	The number of seconds that the credential is capable of accepting security contexts under the specified mechanism. If the credential can no longer be used to accept contexts, or if the credential usage for this mechanism is <code>GSS_C_INITIATE</code> , this parameter will be set to 0. Specify <code>NULL</code> if this parameter is not required.														
<i>cred_usage</i>	How the credential may be used with the specified mechanism. The <i>cred_usage</i> parameter may contain one of the following values: <code>GSS_C_INITIATE</code> , <code>GSS_C_ACCEPT</code> , or <code>GSS_C_BOTH</code> . Specify <code>NULL</code> if this parameter is not required.														
<b>RETURN VALUES</b>	<code>gss_inquire_cred_by_mech()</code> may return the following status codes: <table><tr><td><code>GSS_S_COMPLETE</code></td><td>Successful completion.</td></tr><tr><td><code>GSS_S_NO_CRED</code></td><td>The referenced credentials cannot be accessed.</td></tr><tr><td><code>GSS_S_DEFECTIVE_CREDENTIAL</code></td><td>The referenced credentials are invalid..</td></tr></table>	<code>GSS_S_COMPLETE</code>	Successful completion.	<code>GSS_S_NO_CRED</code>	The referenced credentials cannot be accessed.	<code>GSS_S_DEFECTIVE_CREDENTIAL</code>	The referenced credentials are invalid..								
<code>GSS_S_COMPLETE</code>	Successful completion.														
<code>GSS_S_NO_CRED</code>	The referenced credentials cannot be accessed.														
<code>GSS_S_DEFECTIVE_CREDENTIAL</code>	The referenced credentials are invalid..														

gss\_inquire\_cred\_by\_mech(3GSS)

GSS\_S\_CREDENTIALS\_EXPIRED

The credentials cannot be added because they have expired.

GSS\_S\_FAILURE

The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the *minor\_status* parameter details the error condition.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO**

`gss_release_name(3GSS)`, `attributes(5)`

GSS-API Programming Guide

## gss\_inquire\_mechs\_for\_name(3GSS)

<b>NAME</b>	<code>gss_inquire_mechs_for_name</code> – list mechanisms that support the specified name-type
<b>SYNOPSIS</b>	<pre>cc [flag...] file... -lgss [library...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_inquire_mechs_for_name</b>(OM_uint32 *minor_status, const     gss_name_t input_name, gss_OID_set *mech_types);</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_inquire_mechs_for_name()</code> function returns the set of mechanisms supported by the GSS-API that may be able to process the specified name. Each mechanism returned will recognize at least one element within the internal name.</p> <p>Some implementations of the GSS-API may perform this test by checking nametype information contained within the passed name and registration information provided by individual mechanisms. This means that the <i>mech_types</i> set returned by the function may indicate that a particular mechanism will understand the name, when in fact the mechanism would refuse to accept the name as input to <code>gss_canonicalize_name(3GSS)</code>, <code>gss_init_sec_context(3GSS)</code>, <code>gss_acquire_cred(3GSS)</code>, or <code>gss_add_cred(3GSS)</code>, due to some property of the name itself rather than the name-type. Therefore, this function should be used only as a pre-filter for a call to a subsequent mechanism-specific function.</p>
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_inquire_mechs_for_name()</code> follow in alphabetical order:</p> <p><i>minor_status</i>      Mechanism-specific status code.</p> <p><i>input_name</i>        The name to which the inquiry relates.</p> <p><i>mech_types</i>        Set of mechanisms that may support the specified name. The returned OID set must be freed by the caller after use with a call to <code>gss_release_oid_set(3GSS)</code>.</p>
<b>RETURN VALUES</b>	<p>The <code>gss_inquire_mechs_for_name()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>      Successful completion.</p> <p><code>GSS_S_BAD_NAME</code>        The <i>input_name</i> parameter was ill-formed.</p> <p><code>GSS_S_BAD_NAME_TYPE</code>    The <i>input_name</i> parameter contained an invalid or unsupported type of name.</p> <p><code>GSS_S_FAILURE</code>        The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

gss\_inquire\_mechs\_for\_name(3GSS)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** gss\_acquire\_cred(3GSS), gss\_add\_cred(3GSS),  
gss\_canonicalize\_name(3GSS), gss\_init\_sec\_context(3GSS),  
gss\_release\_oid\_set(3GSS), attributes(5)

GSS-API Programming Guide

## gss\_inquire\_names\_for\_mech(3GSS)

<b>NAME</b>	<code>gss_inquire_names_for_mech</code> – list the name-types supported by the specified mechanism								
<b>SYNOPSIS</b>	<pre>cc [flag...] file... -lgss [library...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_inquire_names_for_mech</b>(OM_uint32 *minor_status, const     gss_OID mechanism, gss_OID_set *name_types);</pre>								
<b>DESCRIPTION</b>	The <code>gss_inquire_names_for_mech()</code> function returns the set of name-types supported by the specified mechanism.								
<b>PARAMETERS</b>	The parameter descriptions for <code>gss_inquire_names_for_mech()</code> follow:  <i>minor_status</i> A mechanism-specific status code.  <i>mechanism</i> The mechanism to be interrogated.  <i>name_types</i> Set of name-types supported by the specified mechanism. The returned OID set must be freed by the application after use with a call to <code>gss_release_oid_set(3GSS)</code> .								
<b>RETURN VALUES</b>	The <code>gss_inquire_names_for_mech()</code> function may return the following values:  GSS_S_COMPLETE    Successful completion.  GSS_S_FAILURE     The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.								
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:  <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWgss (32-bit)</td></tr><tr><td></td><td>SUNWgssx (64-bit)</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<code>gss_release_oid_set(3GSS)</code> , <code>attributes(5)</code>  GSS-API Programming Guide								

<b>NAME</b>	gss_oid_to_str – convert an OID to a string
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  <b>gss_oid_to_str</b>(OM_uint32 *minor_status, const gss_OID *oid,                 gss_buffer_t oid_str);</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_oid_to_str()</code> function converts a GSS-API OID structure to a string. You can use the function to convert the name of a mechanism from an OID to a simple string. This function is a convenience function, as is its complementary function, <code>gss_str_to_oid(3GSS)</code>.</p> <p>If an OID must be created, use <code>gss_create_empty_oid_set(3GSS)</code> and <code>gss_add_oid_set_member()</code> (3GSS) to create it. OIDs created in this way must be released with <code>gss_release_oid_set(3GSS)</code>. However, it is strongly suggested that applications use the default GSS-API mechanism instead of creating an OID for a specific mechanism.</p>
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_oid_to_str()</code> are as follows:</p> <p><i>minor_status</i>            Status code returned by underlying mechanism.</p> <p><i>oid</i>                      GSS-API OID structure to convert.</p> <p><i>oid_str</i>                 String to receive converted OID.</p>
<b>RETURN VALUES</b>	<p><code>gss_oid_to_str()</code> returns one of the following status codes:</p> <p><code>GSS_S_CALL_INACCESSIBLE_READ</code>    A required input parameter could not be read.</p> <p><code>GSS_S_CALL_INACCESSIBLE_WRITE</code>    A required output parameter could not be written.</p> <p><code>GSS_S_COMPLETE</code>                    Successful completion.</p> <p><code>GSS_S_FAILURE</code>                    The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

`gss_oid_to_str(3GSS)`

**SEE ALSO** | `gss_add_oid_set_member()`(3GSS), `gss_create_empty_oid_set`(3GSS),  
`gss_release_oid_set`(3GSS), `gss_str_to_oid`(3GSS), `attributes`(5)

GSS-API Programming Guide

**WARNINGS** | This function is included for compatibility only with programs using earlier versions of the GSS-API and should not be used for new programs. Other implementations of the GSS-API might not support this function, so portable programs should not rely on it. Sun might not continue to support this function.



<b>NAME</b>	gss_process_context_token – pass asynchronous token to security service
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_process_context_token(OM_uint32 *minor_status, const     gss_ctx_id_t context_handle, const gss_buffer_t token_buffer);</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_process_context_token()</code> function provides a way to pass an asynchronous token to the security service. Most context-level tokens are emitted and processed synchronously by <code>gss_init_sec_context()</code> and <code>gss_accept_sec_context()</code>, and the application is informed as to whether further tokens are expected by the <code>GSS_C_CONTINUE_NEEDED</code> major status bit. Occasionally, a mechanism might need to emit a context-level token at a point when the peer entity is not expecting a token. For example, the initiator's final call to <code>gss_init_sec_context()</code> may emit a token and return a status of <code>GSS_S_COMPLETE</code>, but the acceptor's call to <code>gss_accept_sec_context()</code> might fail. The acceptor's mechanism might want to send a token containing an error indication to the initiator, but the initiator is not expecting a token at this point, believing that the context is fully established. <code>gss_process_context_token()</code> provides a way to pass such a token to the mechanism at any time.</p> <p>This function is provided for compatibility with the GSS-API version 1. Because <code>gss_delete_sec_context()</code> no longer returns a valid <i>output_token</i> to be sent to <code>gss_process_context_token()</code>, applications using a newer version of the GSS-API do not need to rely on this function.</p>
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_process_context_token()</code> are as follows:</p> <p><i>minor_status</i>            A mechanism-specific status code.</p> <p><i>context_handle</i>        Context handle of context on which token is to be processed.</p> <p><i>token_buffer</i>            Token to process.</p>
<b>RETURN VALUES</b>	<p><code>gss_process_context_token()</code> returns one of the following status codes:</p> <p><code>GSS_S_COMPLETE</code>            Successful completion.</p> <p><code>GSS_S_DEFECTIVE_TOKEN</code>    Indicates that consistency checks performed on the token failed.</p> <p><code>GSS_S_NO_CONTEXT</code>        The <i>context_handle</i> did not refer to a valid context.</p> <p><code>GSS_S_FAILURE</code>            The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

gss\_process\_context\_token(3GSS)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT Level	Safe

**SEE ALSO** `gss_accept_sec_context(3GSS)`, `gss_delete_sec_context(3GSS)`,  
`gss_init_sec_context(3GSS)`, `attributes(5)`

*GSS-API Programming Guide*

<b>NAME</b>	gss_release_buffer – free buffer storage allocated by a GSS-API function								
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_release_buffer(OM_uint32 *minor_status,                              gss_buffer_tbuffer);</pre>								
<b>DESCRIPTION</b>	The <code>gss_release_buffer()</code> function frees buffer storage allocated by a GSS-API function. The <code>gss_release_buffer()</code> function also zeros the length field in the descriptor to which the buffer parameter refers, while the GSS-API function sets the pointer field in the descriptor to NULL. Any buffer object returned by a GSS-API function may be passed to <code>gss_release_buffer()</code> , even if no storage is associated with the buffer.								
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_release_buffer()</code> follow:</p> <p><i>minor_status</i>      Mechanism-specific status code.</p> <p><i>buffer</i>              The storage associated with the buffer will be deleted. The <code>gss_buffer_desc()</code> object will not be freed; however, its length field will be zeroed.</p>								
<b>RETURN VALUES</b>	<p>The <code>gss_release_buffer()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>      Successful completion</p> <p><code>GSS_S_FAILURE</code>      The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWgss (32-bit)</td> </tr> <tr> <td></td> <td>SUNWgssx (64-bit)</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<p><code>attributes(5)</code></p> <p>GSS-API Programming Guide</p>								

## gss\_release\_cred(3GSS)

<b>NAME</b>	gss_release_cred – discard a credential handle								
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_release_cred(OM_uint32 *minor_status, gss_cred_id_t     *cred_handle) ;</pre>								
<b>DESCRIPTION</b>	The <code>gss_release_cred()</code> function informs the GSS-API that the specified credential handle is no longer required by the application and frees the associated resources. The <code>cred_handle</code> parameter is set to <code>GSS_C_NO_CREDENTIAL</code> when this call completes successfully.								
<b>PARAMETERS</b>	The parameter descriptions for <code>gss_release_cred()</code> follow:  <i>minor_status</i> A mechanism specific status code.  <i>cred_handle</i> An opaque handle that identifies the credential to be released. If <code>GSS_C_NO_CREDENTIAL</code> is specified, the <code>gss_release_cred()</code> function will complete successfully, but it will do nothing.								
<b>RETURN VALUES</b>	<code>gss_release_cred()</code> may return the following status codes:  <code>GSS_S_COMPLETE</code> Successful completion.  <code>GSS_S_NO_CRED</code> The referenced credentials cannot be accessed.  <code>GSS_S_FAILURE</code> The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.								
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:  <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWgss (32-bit)</td></tr><tr><td></td><td>SUNWgssx (64-bit)</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<code>attributes(5)</code>  GSS-API Programming Guide								

<b>NAME</b>	gss_release_name – discard an internal-form name
<b>SYNOPSIS</b>	<pre>cc [flag...] file... -lgss [library...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_release_name(OM_uint32 *minor_status, gss_name_t     *name);</pre>
<b>DESCRIPTION</b>	The <code>gss_release_name()</code> function frees GSS-API-allocated storage associated with an internal-form name. The <code>name</code> is set to <code>GSS_C_NO_NAME</code> on successful completion of this call.
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_release_name()</code> follow:</p> <p><code>minor_status</code>     A mechanism-specific status code.</p> <p><code>name</code>                The name to be deleted.</p>
<b>RETURN VALUES</b>	<p>The <code>gss_release_name()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>     Successful completion.</p> <p><code>GSS_S_BAD_NAME</code>     The <code>name</code> parameter did not contain a valid name.</p> <p><code>GSS_S_FAILURE</code>     The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <code>minor_status</code> parameter details the error condition.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `attributes(5)`  
GSS-API Programming Guide

## gss\_release\_oid(3GSS)

<b>NAME</b>	gss_release_oid – release an object identifier								
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  <b>gss_release_oid</b>(OM_uint32 *minor_status, const gss_OID *oid);</pre>								
<b>DESCRIPTION</b>	<p>The <code>gss_release_oid()</code> function deletes an OID. Such an OID might have been created with <code>gss_str_to_oid()</code>.</p> <p>Since creating and deleting individual OIDs is discouraged, it is preferable to use <code>gss_release_oid_set()</code> if it is necessary to deallocate a set of OIDs.</p>								
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_release_oid()</code> are as follows:</p> <p><i>minor_status</i>      A mechanism-specific status code.</p> <p><i>oid</i>                The object identifier of the mechanism to be deleted.</p>								
<b>RETURN VALUES</b>	<p><code>gss_release_oid()</code> returns one of the following status codes:</p> <p><code>GSS_S_COMPLETE</code>    Successful completion.</p> <p><code>GSS_S_FAILURE</code>     The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWgss (32-bit)</td></tr><tr><td></td><td>SUNWgssx (64-bit)</td></tr><tr><td>MT Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT Level	Safe								
<b>SEE ALSO</b>	<p><code>gss_release_oid_set(3GSS)</code>, <code>gss_str_to_oid(3GSS)</code>, <code>attributes(5)</code></p> <p>GSS-API Programming Guide</p>								
<b>WARNINGS</b>	<p>This function is included for compatibility only with programs using earlier versions of the GSS-API and should not be used for new programs. Other implementations of the GSS-API might not support this function, so portable programs should not rely on it. Sun might not continue to support this function.</p>								

<b>NAME</b>	<code>gss_release_oid_set</code> – free storage associated with a GSS-API-generated <code>gss_OID_set</code> object
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_release_oid_set</b>(OM_uint32 *minor_status, gss_OID_set     *set);</pre>
<b>DESCRIPTION</b>	<p>The <code>gss_release_oid_set()</code> function frees storage associated with a GSS-API-generated <code>gss_OID_set</code> object. The <code>set</code> parameter must refer to an OID-set that was returned from a GSS-API function. The <code>gss_release_oid_set()</code> function will free the storage associated with each individual member OID, the OID <code>set</code>'s elements array, and <code>gss_OID_set_desc</code>.</p> <p><code>gss_OID_set</code> is set to <code>GSS_C_NO_OID_SET</code> on successful completion of this function.</p>
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_release_oid_set()</code> follow:</p> <p><i>minor_status</i>            A mechanism-specific status code</p> <p><i>set</i>                      Storage associated with the <code>gss_OID_set</code> will be deleted</p>
<b>RETURN VALUES</b>	<p>The <code>gss_release_oid_set()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>        Successful completion</p> <p><code>GSS_S_FAILURE</code>        The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `attributes(5)`  
GSS-API Programming Guide

## gss\_str\_to\_oid(3GSS)

<b>NAME</b>	gss_str_to_oid – convert a string to an OID								
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_str_to_oid(OM_uint32 *minor_status, const gss_buffer_t oid_str, gss_OID *oid);</pre>								
<b>DESCRIPTION</b>	<p>The <code>gss_str_to_oid()</code> function converts a string to a GSS-API OID structure. You can use the function to convert a simple string to an OID to . This function is a convenience function, as is its complementary function, <code>gss_oid_to_str(3GSS)</code>.</p> <p>OIDs created with <code>gss_str_to_oid()</code> must be deallocated through <code>gss_release_oid(3GSS)</code>, if available. If an OID must be created, use <code>gss_create_empty_oid_set(3GSS)</code> and <code>gss_add_oid_set_member()</code> (3GSS) to create it. OIDs created in this way must be released with <code>gss_release_oid_set(3GSS)</code>. However, it is strongly suggested that applications use the default GSS-API mechanism instead of creating an OID for a specific mechanism.</p>								
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_str_to_oid()</code> are as follows:</p> <p><i>minor_status</i>            Status code returned by underlying mechanism.</p> <p><i>oid</i>                        GSS-API OID structure to receive converted string.</p> <p><i>oid_str</i>                    String to convert.</p>								
<b>RETURN VALUES</b>	<p><code>gss_str_to_oid()</code> returns one of the following status codes:</p> <table border="0"> <tr> <td style="vertical-align: top;">GSS_S_CALL_INACCESSIBLE_READ</td> <td>A required input parameter could not be read.</td> </tr> <tr> <td style="vertical-align: top;">GSS_S_CALL_INACCESSIBLE_WRITE</td> <td>A required output parameter could not be written.</td> </tr> <tr> <td style="vertical-align: top;">GSS_S_COMPLETE</td> <td>Successful completion.</td> </tr> <tr> <td style="vertical-align: top;">GSS_S_FAILURE</td> <td>The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</td> </tr> </table>	GSS_S_CALL_INACCESSIBLE_READ	A required input parameter could not be read.	GSS_S_CALL_INACCESSIBLE_WRITE	A required output parameter could not be written.	GSS_S_COMPLETE	Successful completion.	GSS_S_FAILURE	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.
GSS_S_CALL_INACCESSIBLE_READ	A required input parameter could not be read.								
GSS_S_CALL_INACCESSIBLE_WRITE	A required output parameter could not be written.								
GSS_S_COMPLETE	Successful completion.								
GSS_S_FAILURE	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWgss (32-bit)</td> </tr> <tr> <td></td> <td>SUNWgssx (64-bit)</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)		
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								



gss\_str\_to\_oid(3GSS)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `gss_add_oid_set_member()`(3GSS), `gss_create_empty_oid_set`(3GSS), `gss_oid_to_str`(3GSS), `gss_release_oid_set`(3GSS), `attributes`(5)

GSS-API Programming Guide

**WARNINGS** This function is included for compatibility only with programs using earlier versions of the GSS-API and should not be used for new programs. Other implementations of the GSS-API might not support this function, so portable programs should not rely on it. Sun might not continue to support this function.

## gss\_test\_oid\_set\_member(3GSS)

<b>NAME</b>	gss_test_oid_set_member – interrogate an object identifier set								
<b>SYNOPSIS</b>	<pre>cc -flag ... file ...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 gss_test_oid_set_member(OM_uint32 *minor_status, const     gss_OID member, const gss_OID_set set, int *present);</pre>								
<b>DESCRIPTION</b>	The <code>gss_test_oid_set_member()</code> function interrogates an object identifier set to determine if a specified object identifier is a member. This function should be used with OID sets returned by <code>gss_indicate_mechs(3GSS)</code> , <code>gss_acquire_cred(3GSS)</code> , and <code>gss_inquire_cred(3GSS)</code> , but it will also work with user-generated sets.								
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_test_oid_set_member()</code> follow:</p> <p><i>minor_status</i>     A mechanism-specific status code</p> <p><i>member</i>            An object identifier whose presence is to be tested</p> <p><i>set</i>                An object identifier set.</p> <p><i>present</i>            The value of <i>present</i> is non-zero if the specified OID is a member of the set; if not, the value of <i>present</i> is zero.</p>								
<b>RETURN VALUES</b>	<p>The <code>gss_test_oid_set_member()</code> function may return the following status codes:</p> <p><code>GSS_S_COMPLETE</code>     Successful completion</p> <p><code>GSS_S_FAILURE</code>     The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.</p>								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWgss (32-bit)</td> </tr> <tr> <td></td> <td>SUNWgssx (64-bit)</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWgss (32-bit)		SUNWgssx (64-bit)	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWgss (32-bit)								
	SUNWgssx (64-bit)								
MT-Level	Safe								
<b>SEE ALSO</b>	<p><code>gss_acquire_cred(3GSS)</code>, <code>gss_indicate_mechs(3GSS)</code>, <code>gss_inquire_cred(3GSS)</code>, <code>attributes(5)</code></p> <p>GSS-API Programming Guide</p>								

<b>NAME</b>	gss_wrap – verify a message with attached cryptographic message												
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_unwrap</b>(OM_uint32 *minor_status, const gss_ctx_id_t     context_handle, const gss_buffer_t input_message_buffer, gss_buffer_t     output_message_buffer, int *conf_state, gss_qop_t *qop_state);</pre>												
<b>DESCRIPTION</b>	<p>The <code>gss_unwrap()</code> function converts a message previously protected by <code>gss_wrap(3GSS)</code> back to a usable form, verifying the embedded MIC. The <code>conf_state</code> parameter indicates whether the message was encrypted; the <code>qop_state</code> parameter indicates the strength of protection that was used to provide the confidentiality and integrity services.</p> <p>Since some application-level protocols may wish to use tokens emitted by <code>gss_wrap(3GSS)</code> to provide secure framing, the GSS-API supports the wrapping and unwrapping of zero-length messages.</p>												
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_unwrap()</code> follow:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><i>minor_status</i></td> <td>The status code returned by the underlying mechanism.</td> </tr> <tr> <td><i>context_handle</i></td> <td>Identifies the context on which the message arrived.</td> </tr> <tr> <td><i>input_message_buffer</i></td> <td>The message to be protected.</td> </tr> <tr> <td><i>output_message_buffer</i></td> <td>The buffer to receive the unwrapped message. Storage associated with this buffer must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code>.</td> </tr> <tr> <td><i>conf_state</i></td> <td>If the value of <i>conf_state</i> is non-zero, then confidentiality and integrity protection were used. If the value is zero, only integrity service was used. Specify NULL if this parameter is not required.</td> </tr> <tr> <td><i>qop_state</i></td> <td>Specifies the quality of protection provided. Specify NULL if this parameter is not required.</td> </tr> </table>	<i>minor_status</i>	The status code returned by the underlying mechanism.	<i>context_handle</i>	Identifies the context on which the message arrived.	<i>input_message_buffer</i>	The message to be protected.	<i>output_message_buffer</i>	The buffer to receive the unwrapped message. Storage associated with this buffer must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code> .	<i>conf_state</i>	If the value of <i>conf_state</i> is non-zero, then confidentiality and integrity protection were used. If the value is zero, only integrity service was used. Specify NULL if this parameter is not required.	<i>qop_state</i>	Specifies the quality of protection provided. Specify NULL if this parameter is not required.
<i>minor_status</i>	The status code returned by the underlying mechanism.												
<i>context_handle</i>	Identifies the context on which the message arrived.												
<i>input_message_buffer</i>	The message to be protected.												
<i>output_message_buffer</i>	The buffer to receive the unwrapped message. Storage associated with this buffer must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code> .												
<i>conf_state</i>	If the value of <i>conf_state</i> is non-zero, then confidentiality and integrity protection were used. If the value is zero, only integrity service was used. Specify NULL if this parameter is not required.												
<i>qop_state</i>	Specifies the quality of protection provided. Specify NULL if this parameter is not required.												
<b>RETURN VALUES</b>	<p><code>gss_unwrap()</code> may return the following status codes:</p> <table border="0"> <tr> <td style="padding-right: 20px;">GSS_S_COMPLETE</td> <td>Successful completion.</td> </tr> <tr> <td>GSS_S_DEFECTIVE_TOKEN</td> <td>The token failed consistency checks.</td> </tr> <tr> <td>GSS_S_BAD_SIG</td> <td>The MIC was incorrect.</td> </tr> <tr> <td>GSS_S_DUPLICATE_TOKEN</td> <td>The token was valid, and contained a correct MIC for the message, but it had already been processed.</td> </tr> <tr> <td>GSS_S_OLD_TOKEN</td> <td>The token was valid, and contained a correct MIC for the message, but it is too old to check for duplication.</td> </tr> </table>	GSS_S_COMPLETE	Successful completion.	GSS_S_DEFECTIVE_TOKEN	The token failed consistency checks.	GSS_S_BAD_SIG	The MIC was incorrect.	GSS_S_DUPLICATE_TOKEN	The token was valid, and contained a correct MIC for the message, but it had already been processed.	GSS_S_OLD_TOKEN	The token was valid, and contained a correct MIC for the message, but it is too old to check for duplication.		
GSS_S_COMPLETE	Successful completion.												
GSS_S_DEFECTIVE_TOKEN	The token failed consistency checks.												
GSS_S_BAD_SIG	The MIC was incorrect.												
GSS_S_DUPLICATE_TOKEN	The token was valid, and contained a correct MIC for the message, but it had already been processed.												
GSS_S_OLD_TOKEN	The token was valid, and contained a correct MIC for the message, but it is too old to check for duplication.												

## gss\_unwrap(3GSS)

GSS_S_UNSEQ_TOKEN	The token was valid, and contained a correct MIC for the message, but has been verified out of sequence; a later token has already been received.
GSS_S_GAP_TOKEN	The token was valid, and contained a correct MIC for the message, but has been verified out of sequence; an earlier expected token has not yet been received.
GSS_S_CONTEXT_EXPIRED	The context has already expired.
GSS_S_NO_CONTEXT	The <i>context_handle</i> parameter did not identify a valid context.
GSS_S_FAILURE	The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the <i>minor_status</i> parameter details the error condition.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_release_buffer(3GSS)`, `gss_wrap(3GSS)`, `attributes(5)`

GSS-API Programming Guide

<b>NAME</b>	gss_verify_mic – verify integrity of a received message														
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_verify_mic</b>(OM_uint32 *minor_status, const gss_ctx_id_t     context_handle, const gss_buffer_t message_buffer, const     gss_buffer_t token_buffer, gss_qop_t *qop_state);</pre>														
<b>DESCRIPTION</b>	<p>The <code>gss_verify_mic()</code> function verifies that a cryptographic MIC, contained in the token parameter, fits the supplied message. The <code>qop_state</code> parameter allows a message recipient to determine the strength of protection that was applied to the message.</p> <p>Since some application-level protocols may wish to use tokens emitted by <code>gss_wrap(3GSS)</code> to provide secure framing, the GSS-API supports the calculation and verification of MICs over zero-length messages.</p>														
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_verify_mic()</code> follow:</p> <p><i>minor_status</i>      The status code returned by the underlying mechanism.</p> <p><i>context_handle</i>    Identifies the context on which the message arrived.</p> <p><i>message_buffer</i>    The message to be verified.</p> <p><i>token_buffer</i>      The token associated with the message.</p> <p><i>qop_state</i>         Specifies the quality of protection gained from the MIC. Specify NULL if this parameter is not required.</p>														
<b>RETURN VALUES</b>	<p><code>gss_verify_mic()</code> may return the following status codes:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>GSS_S_COMPLETE</code></td> <td>Successful completion.</td> </tr> <tr> <td><code>GSS_S_DEFECTIVE_TOKEN</code></td> <td>The token failed consistency checks.</td> </tr> <tr> <td><code>GSS_S_BAD_SIG</code></td> <td>The MIC was incorrect.</td> </tr> <tr> <td><code>GSS_S_DUPLICATE_TOKEN</code></td> <td>The token was valid and contained a correct MIC for the message, but it had already been processed.</td> </tr> <tr> <td><code>GSS_S_OLD_TOKEN</code></td> <td>The token was valid and contained a correct MIC for the message, but it is too old to check for duplication.</td> </tr> <tr> <td><code>GSS_S_UNSEQ_TOKEN</code></td> <td>The token was valid and contained a correct MIC for the message, but it has been verified out of sequence; a later token has already been received.</td> </tr> <tr> <td><code>GSS_S_GAP_TOKEN</code></td> <td>The token was valid and contained a correct MIC for the message, but it has been verified out of sequence; an earlier expected token has not yet been received.</td> </tr> </table>	<code>GSS_S_COMPLETE</code>	Successful completion.	<code>GSS_S_DEFECTIVE_TOKEN</code>	The token failed consistency checks.	<code>GSS_S_BAD_SIG</code>	The MIC was incorrect.	<code>GSS_S_DUPLICATE_TOKEN</code>	The token was valid and contained a correct MIC for the message, but it had already been processed.	<code>GSS_S_OLD_TOKEN</code>	The token was valid and contained a correct MIC for the message, but it is too old to check for duplication.	<code>GSS_S_UNSEQ_TOKEN</code>	The token was valid and contained a correct MIC for the message, but it has been verified out of sequence; a later token has already been received.	<code>GSS_S_GAP_TOKEN</code>	The token was valid and contained a correct MIC for the message, but it has been verified out of sequence; an earlier expected token has not yet been received.
<code>GSS_S_COMPLETE</code>	Successful completion.														
<code>GSS_S_DEFECTIVE_TOKEN</code>	The token failed consistency checks.														
<code>GSS_S_BAD_SIG</code>	The MIC was incorrect.														
<code>GSS_S_DUPLICATE_TOKEN</code>	The token was valid and contained a correct MIC for the message, but it had already been processed.														
<code>GSS_S_OLD_TOKEN</code>	The token was valid and contained a correct MIC for the message, but it is too old to check for duplication.														
<code>GSS_S_UNSEQ_TOKEN</code>	The token was valid and contained a correct MIC for the message, but it has been verified out of sequence; a later token has already been received.														
<code>GSS_S_GAP_TOKEN</code>	The token was valid and contained a correct MIC for the message, but it has been verified out of sequence; an earlier expected token has not yet been received.														

## gss\_verify\_mic(3GSS)

GSS\_S\_CONTEXT\_EXPIRED

The context has already expired.

GSS\_S\_NO\_CONTEXT

The *context\_handle* parameter did not identify a valid context.

GSS\_S\_FAILURE

The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the *minor\_status* parameter details the error condition.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_wrap(3GSS)`, `attributes(5)`

GSS-API Programming Guide

<b>NAME</b>	gss_wrap – attach a cryptographic message														
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_wrap</b>(OM_uint32 *minor_status, const gss_ctx_id_t     context_handle, int conf_req_flag, gss_qop_t qop_req, const     gss_buffer_t input_message_buffer, int *conf_state, gss_buffer_t     output_message_buffer);</pre>														
<b>DESCRIPTION</b>	<p>The <code>gss_wrap()</code> function attaches a cryptographic MIC and optionally encrypts the specified <code>input_message</code>. The <code>output_message</code> contains both the MIC and the message. The <code>qop_req</code> parameter allows a choice between several cryptographic algorithms, if supported by the chosen mechanism.</p> <p>Since some application-level protocols may wish to use tokens emitted by <code>gss_wrap()</code> to provide secure framing, the GSS-API supports the wrapping of zero-length messages.</p>														
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_wrap()</code> follow:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><i>minor_status</i></td> <td>The status code returned by the underlying mechanism.</td> </tr> <tr> <td><i>context_handle</i></td> <td>Identifies the context on which the message will be sent.</td> </tr> <tr> <td><i>conf_req_flag</i></td> <td>If the value of <i>conf_req_flag</i> is non-zero, both confidentiality and integrity services are requested. If the value is zero, then only integrity service is requested.</td> </tr> <tr> <td><i>qop_req</i></td> <td>Specifies the required quality of protection. A mechanism-specific default may be requested by setting <i>qop_req</i> to <code>GSS_C_QOP_DEFAULT</code>. If an unsupported protection strength is requested, <code>gss_wrap()</code> will return a <i>major_status</i> of <code>GSS_S_BAD_QOP</code>.</td> </tr> <tr> <td><i>input_message_buffer</i></td> <td>The message to be protected.</td> </tr> <tr> <td><i>conf_state</i></td> <td>If the value of <i>conf_state</i> is non-zero, confidentiality, data origin authentication, and integrity services have been applied. If the value is zero, then integrity services have been applied. Specify <code>NULL</code> if this parameter is not required.</td> </tr> <tr> <td><i>output_message_buffer</i></td> <td>The buffer to receive the protected message. Storage associated with this message must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code>.</td> </tr> </table>	<i>minor_status</i>	The status code returned by the underlying mechanism.	<i>context_handle</i>	Identifies the context on which the message will be sent.	<i>conf_req_flag</i>	If the value of <i>conf_req_flag</i> is non-zero, both confidentiality and integrity services are requested. If the value is zero, then only integrity service is requested.	<i>qop_req</i>	Specifies the required quality of protection. A mechanism-specific default may be requested by setting <i>qop_req</i> to <code>GSS_C_QOP_DEFAULT</code> . If an unsupported protection strength is requested, <code>gss_wrap()</code> will return a <i>major_status</i> of <code>GSS_S_BAD_QOP</code> .	<i>input_message_buffer</i>	The message to be protected.	<i>conf_state</i>	If the value of <i>conf_state</i> is non-zero, confidentiality, data origin authentication, and integrity services have been applied. If the value is zero, then integrity services have been applied. Specify <code>NULL</code> if this parameter is not required.	<i>output_message_buffer</i>	The buffer to receive the protected message. Storage associated with this message must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code> .
<i>minor_status</i>	The status code returned by the underlying mechanism.														
<i>context_handle</i>	Identifies the context on which the message will be sent.														
<i>conf_req_flag</i>	If the value of <i>conf_req_flag</i> is non-zero, both confidentiality and integrity services are requested. If the value is zero, then only integrity service is requested.														
<i>qop_req</i>	Specifies the required quality of protection. A mechanism-specific default may be requested by setting <i>qop_req</i> to <code>GSS_C_QOP_DEFAULT</code> . If an unsupported protection strength is requested, <code>gss_wrap()</code> will return a <i>major_status</i> of <code>GSS_S_BAD_QOP</code> .														
<i>input_message_buffer</i>	The message to be protected.														
<i>conf_state</i>	If the value of <i>conf_state</i> is non-zero, confidentiality, data origin authentication, and integrity services have been applied. If the value is zero, then integrity services have been applied. Specify <code>NULL</code> if this parameter is not required.														
<i>output_message_buffer</i>	The buffer to receive the protected message. Storage associated with this message must be freed by the application after use with a call to <code>gss_release_buffer(3GSS)</code> .														
<b>RETURN VALUES</b>	<p><code>gss_wrap()</code> may return the following status codes:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>GSS_S_COMPLETE</code></td> <td>Successful completion.</td> </tr> <tr> <td><code>GSS_S_CONTEXT_EXPIRED</code></td> <td>The context has already expired.</td> </tr> </table>	<code>GSS_S_COMPLETE</code>	Successful completion.	<code>GSS_S_CONTEXT_EXPIRED</code>	The context has already expired.										
<code>GSS_S_COMPLETE</code>	Successful completion.														
<code>GSS_S_CONTEXT_EXPIRED</code>	The context has already expired.														

## gss\_wrap(3GSS)

GSS\_S\_NO\_CONTEXT

The *context\_handle* parameter did not identify a valid context.

GSS\_S\_BAD\_QOP

The specified QOP is not supported by the mechanism.

GSS\_S\_FAILURE

The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the *minor\_status* parameter details the error condition.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT-Level	Safe

**SEE ALSO** `gss_release_buffer(3GSS)`, `attributes(5)`

GSS-API Programming Guide



`gss_wrap_size_limit(3GSS)`

<b>NAME</b>	<code>gss_wrap_size_limit</code> – allow application to determine maximum message size with resulting output token of a specified maximum size												
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lgss [library ...] #include &lt;gssapi/gssapi.h&gt;  OM_uint32 <b>gss_process_context_token</b>(OM_uint32 *minor_status, const     gss_ctx_id_t context_handle, int conf_req_flag, gss_qop_t qop_req,     OM_uint32 req_output_size, OM_uint32 *max_input_size);</pre>												
<b>DESCRIPTION</b>	<p>The <code>gss_wrap_size_limit()</code> function allows an application to determine the maximum message size that, if presented to <code>gss_wrap()</code> with the same <code>conf_req_flag</code> and <code>qop_req</code> parameters, results in an output token containing no more than <code>req_output_size</code> bytes. This call is intended for use by applications that communicate over protocols that impose a maximum message size. It enables the application to fragment messages prior to applying protection. The GSS-API detects invalid QOP values when <code>gss_wrap_size_limit()</code> is called. This routine guarantees only a maximum message size, not the availability of specific QOP values for message protection.</p> <p>Successful completion of <code>gss_wrap_size_limit()</code> does not guarantee that <code>gss_wrap()</code> will be able to protect a message of length <code>max_input_size</code> bytes, since this ability might depend on the availability of system resources at the time that <code>gss_wrap()</code> is called.</p>												
<b>PARAMETERS</b>	<p>The parameter descriptions for <code>gss_wrap_size_limit()</code> are as follows:</p> <table><tr><td><code>minor_status</code></td><td>A mechanism-specific status code.</td></tr><tr><td><code>context_handle</code></td><td>A handle that refers to the security over which the messages will be sent.</td></tr><tr><td><code>conf_req_flag</code></td><td>Indicates whether <code>gss_wrap()</code> will be asked to apply confidential protection in addition to integrity protection. See <code>gss_wrap(3GSS)</code> for more details.</td></tr><tr><td><code>qop_req</code></td><td>Indicates the level of protection that <code>gss_wrap()</code> will be asked to provide. See <code>gss_wrap(3GSS)</code> for more details.</td></tr><tr><td><code>req_output_size</code></td><td>The desired maximum size for tokens emitted by <code>gss_wrap()</code>.</td></tr><tr><td><code>max_input_size</code></td><td>The maximum input message size that can be presented to <code>gss_wrap()</code> to guarantee that the emitted token will be no larger than <code>req_output_size</code> bytes.</td></tr></table>	<code>minor_status</code>	A mechanism-specific status code.	<code>context_handle</code>	A handle that refers to the security over which the messages will be sent.	<code>conf_req_flag</code>	Indicates whether <code>gss_wrap()</code> will be asked to apply confidential protection in addition to integrity protection. See <code>gss_wrap(3GSS)</code> for more details.	<code>qop_req</code>	Indicates the level of protection that <code>gss_wrap()</code> will be asked to provide. See <code>gss_wrap(3GSS)</code> for more details.	<code>req_output_size</code>	The desired maximum size for tokens emitted by <code>gss_wrap()</code> .	<code>max_input_size</code>	The maximum input message size that can be presented to <code>gss_wrap()</code> to guarantee that the emitted token will be no larger than <code>req_output_size</code> bytes.
<code>minor_status</code>	A mechanism-specific status code.												
<code>context_handle</code>	A handle that refers to the security over which the messages will be sent.												
<code>conf_req_flag</code>	Indicates whether <code>gss_wrap()</code> will be asked to apply confidential protection in addition to integrity protection. See <code>gss_wrap(3GSS)</code> for more details.												
<code>qop_req</code>	Indicates the level of protection that <code>gss_wrap()</code> will be asked to provide. See <code>gss_wrap(3GSS)</code> for more details.												
<code>req_output_size</code>	The desired maximum size for tokens emitted by <code>gss_wrap()</code> .												
<code>max_input_size</code>	The maximum input message size that can be presented to <code>gss_wrap()</code> to guarantee that the emitted token will be no larger than <code>req_output_size</code> bytes.												
<b>RETURN VALUES</b>	<p><code>gss_wrap_size_limit()</code> returns one of the following status codes:</p> <table><tr><td><code>GSS_S_COMPLETE</code></td><td>Successful completion.</td></tr><tr><td><code>GSS_S_NO_CONTEXT</code></td><td>The referenced context could not be accessed.</td></tr><tr><td><code>GSS_S_CONTEXT_EXPIRED</code></td><td>The context has expired.</td></tr></table>	<code>GSS_S_COMPLETE</code>	Successful completion.	<code>GSS_S_NO_CONTEXT</code>	The referenced context could not be accessed.	<code>GSS_S_CONTEXT_EXPIRED</code>	The context has expired.						
<code>GSS_S_COMPLETE</code>	Successful completion.												
<code>GSS_S_NO_CONTEXT</code>	The referenced context could not be accessed.												
<code>GSS_S_CONTEXT_EXPIRED</code>	The context has expired.												

## gss\_wrap\_size\_limit(3GSS)

GSS\_S\_BAD\_QOP

The specified QOP is not supported by the mechanism.

GSS\_S\_FAILURE

The underlying mechanism detected an error for which no specific GSS status code is defined. The mechanism-specific status code reported by means of the *minor\_status* parameter details the error condition.

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWgss (32-bit)
	SUNWgssx (64-bit)
MT Level	Safe

### SEE ALSO

`gss_wrap(3GSS)`, `attributes(5)`

GSS-API Programming Guide

<b>NAME</b>	htonl, htons, ntohl, ntohs – convert values between host and network byte order				
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;arpa/inet.h&gt;  uint32_t <b>htonl</b> (uint32_t <i>hostlong</i>) ; uint16_t <b>htons</b> (uint16_t <i>hostshort</i>) ; uint32_t <b>ntohl</b> (uint32_t <i>netlong</i>) ; uint16_t <b>ntohs</b> (uint16_t <i>netshort</i>) ;</pre>				
<b>DESCRIPTION</b>	<p>These functions convert 16-bit and 32-bit quantities between network byte order and host byte order.</p> <p>The <code>uint32_t</code> and <code>uint16_t</code> types are made available by inclusion of <code>&lt;inttypes.h&gt;</code>.</p>				
<b>USAGE</b>	<p>These functions are most often used in conjunction with Internet addresses and ports as returned by <code>gethostent(3XNET)</code> and <code>getservent(3XNET)</code>.</p> <p>On some architectures these functions are defined as macros that expand to the value of their argument.</p>				
<b>RETURN VALUES</b>	<p>The <code>htonl()</code> and <code>htons()</code> functions return the argument value converted from host to network byte order.</p> <p>The <code>ntohl()</code> and <code>ntohs()</code> functions return the argument value converted from network to host byte order.</p>				
<b>ERRORS</b>	No errors are defined.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>endhostent(3XNET)</code> , <code>endservent(3XNET)</code> , <code>attributes(5)</code>				

## if\_nametoindex(3NSL)

<b>NAME</b>	if_nametoindex, if_indextoname, if_nameindex, if_freenameindex – routines to map Internet Protocol network interface names and interface indexes
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;net/if.h&gt;  unsigned int if_nametoindex(const char *ifname); char *if_indextoname(unsigned int ifindex, char *ifname); struct if_nameindex *if_nameindex(void); void if_freenameindex(struct if_nameindex *ptr);</pre>
<b>DESCRIPTION</b>	<p>This API defines two functions that map between an Internet Protocol network interface name and index, a third function that returns all the interface names and indexes, and a fourth function to return the dynamic memory allocated by the previous function.</p> <p>Network interfaces are normally known by names such as "le0", "sl1", "ppp2", and the like. The <i>ifname</i> argument must point to a buffer of at least IF_NAMESIZE bytes into which the interface name corresponding to the specified index is returned. IF_NAMESIZE is defined in &lt;net/if.h&gt; and its value includes a terminating null byte at the end of the interface name.</p> <p><b>if_nametoindex()</b> The if_nametoindex() function returns the interface index corresponding to the interface name pointed to by the <i>ifname</i> pointer. If the specified interface name does not exist, the return value is 0, and <i>errno</i> is set to ENXIO. If there was a system error, such as running out of memory, the return value is 0 and <i>errno</i> is set to the proper value, for example, ENOMEM.</p> <p><b>if_indextoname()</b> The if_indextoname() function maps an interface index into its corresponding name. This pointer is also the return value of the function. If there is no interface corresponding to the specified index, NULL is returned, and <i>errno</i> is set to ENXIO, if there was a system error, such as running out of memory, if_indextoname() returns NULL and <i>errno</i> would be set to the proper value, for example, ENOMEM.</p> <p><b>*if_nameindex()</b> The if_nameindex() function returns an array of if_nameindex structures, one structure per interface. The if_nameindex structure holds the information about a single interface and is defined when the &lt;net/if.h&gt; header is included:</p> <pre>struct if_nameindex {     unsigned int    if_index; /* 1, 2, ... */     char           *if_name; /* null terminated name: "le0", ... */ };</pre>

## if\_nametoindex(3NSL)

The end of the array of structures is indicated by a structure with an `if_index` of 0 and an `if_name` of NULL. The function returns a null pointer upon an error and sets `errno` to the appropriate value. The memory used for this array of structures along with the interface names pointed to by the `if_name` members is obtained dynamically. This memory is freed by the `if_freenameindex()` function.

### `if_freenameindex()`

The `if_freenameindex()` function frees the dynamic memory that was allocated by `if_nameindex()`. The argument to this function must be a pointer that was returned by `if_nameindex()`.

### PARAMETERS

*ifname* interface name.  
*ifindex* interface index.  
*ptr* pointer returned by `if_nameindex()`.

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
MT Level	MT Safe
Interface Stability	Standard

### SEE ALSO

`ifconfig(1M)`, `attributes(5)`, `if(7P)`

## if\_nametoindex(3XNET)

<b>NAME</b>	if_nametoindex, if_indextoname, if_nameindex, if_freenameindex – functions to map Internet Protocol network interface names and interface indexes
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;net/if.h&gt;  unsigned int if_nametoindex(const char *ifname); char *if_indextoname(unsigned int ifindex, char *ifname); struct if_nameindex *if_nameindex(void); void if_freenameindex(struct if_nameindex *ptr);</pre>
<b>DESCRIPTION</b>	<p>This API defines two functions that map between an Internet Protocol network interface name and index, a third function that returns all the interface names and indexes, and a fourth function to return the dynamic memory allocated by the previous function.</p> <p>Network interfaces are normally known by names such as "le0", "sl1", "ppp2", and the like. The <i>ifname</i> argument must point to a buffer of at least IF_NAMESIZE bytes into which the interface name corresponding to the specified index is returned. IF_NAMESIZE is defined in &lt;net/if.h&gt; and its value includes a terminating null byte at the end of the interface name.</p> <p><b>if_nametoindex()</b> The if_nametoindex() function returns the interface index corresponding to the interface name pointed to by the <i>ifname</i> pointer. If the specified interface name does not exist, the return value is 0, and <i>errno</i> is set to ENXIO. If there was a system error, such as running out of memory, the return value is 0 and <i>errno</i> is set to the proper value, for example, ENOMEM.</p> <p><b>if_indextoname()</b> The if_indextoname() function maps an interface index into its corresponding name. This pointer is also the return value of the function. If there is no interface corresponding to the specified index, NULL is returned, and <i>errno</i> is set to ENXIO, if there was a system error, such as running out of memory, if_indextoname() returns NULL and <i>errno</i> would be set to the proper value, for example, ENOMEM.</p> <p><b>*if_nameindex()</b> The if_nameindex() function returns an array of if_nameindex structures, one structure per interface. The if_nameindex structure holds the information about a single interface and is defined when the &lt;net/if.h&gt; header is included:</p> <pre>struct if_nameindex {     unsigned int  if_index; /* 1, 2, ... */     char          *if_name; /* null terminated name: "le0", ... */ };</pre>

## if\_nametoindex(3XNET)

The end of the array of structures is indicated by a structure with an `if_index` of 0 and an `if_name` of NULL. The function returns a null pointer upon an error and sets `errno` to the appropriate value. The memory used for this array of structures along with the interface names pointed to by the `if_name` members is obtained dynamically. This memory is freed by the `if_freenameindex()` function.

### `if_freenameindex()`

The `if_freenameindex()` function frees the dynamic memory that was allocated by `if_nameindex()`. The argument to this function must be a pointer that was returned by `if_nameindex()`.

### PARAMETERS

*ifname* interface name.  
*ifindex* interface index.  
*ptr* pointer returned by `if_nameindex()`.

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
MT Level	MT Safe
Interface Stability	Standard

### SEE ALSO

`ifconfig(1M)`, `attributes(5)`, `if(7P)`

## inet(3SOCKET)

<b>NAME</b>	inet, inet6, inet_ntop, inet_pton, inet_addr, inet_network, inet_makeaddr, inet_lnaof, inet_netof, inet_ntoa – Internet address manipulation
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt; #include &lt;netinet/in.h&gt; #include &lt;arpa/inet.h&gt;  const char *inet_ntop(int af, const void *addr, char *cp, size_t     size);  int inet_pton(int af, const char *cp, void *addr);  in_addr_t inet_addr(const char *cp);  in_addr_t inet_network(const char *cp);  struct in_addr inet_makeaddr(const int net, const int lna);  int inet_lnaof(const struct in_addr in);  int inet_netof(const struct in_addr in);  char *inet_ntoa(const struct in_addr in);</pre>
<b>DESCRIPTION</b>	<p>The <code>inet_ntop()</code> and <code>inet_pton()</code> routines can manipulate both IPv4 and IPv6 addresses, whereas <code>inet_addr()</code>, <code>inet_network()</code>, <code>inet_makeaddr()</code>, <code>inet_lnaof()</code>, <code>inet_netof()</code>, and <code>inet_ntoa()</code> can only manipulate IPv4 addresses.</p> <p>The <code>inet_ntop()</code> routine converts a numeric address into a string suitable for presentation. The <code>af</code> argument specifies the family of the address. This can be <code>AF_INET</code> or <code>AF_INET6</code>. The <code>addr</code> argument points to a buffer holding an IPv4 address if the <code>af</code> argument is <code>AF_INET</code>, or an IPv6 address if the <code>af</code> argument is <code>AF_INET6</code>; the address must be in network byte order. The <code>cp</code> argument points to a buffer where the routine will store the resulting string. The <code>size</code> argument specifies the size of this buffer. The application must specify a non-NULL <code>cp</code> argument. For IPv6 addresses, the buffer must be at least 46-octets. For IPv4 addresses, the buffer must be at least 16-octets. 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 <code>&lt;netinet/in.h&gt;</code>:</p> <pre>#define INET_ADDRSTRLEN    16 #define INET6_ADDRSTRLEN  46</pre> <p>The <code>inet_pton()</code> routine converts an address in its standard text presentation form into its numeric binary form. The <code>af</code> argument specifies the family of the address. Currently the <code>AF_INET</code> and <code>AF_INET6</code> address families are supported. The <code>cp</code> argument points to the string being passed in. The <code>addr</code> argument points to a buffer into which the routine stores the numeric address. The calling application must ensure that the buffer referred to by <code>addr</code> is large enough to hold the numeric address, at least 4 bytes for <code>AF_INET</code> or 16 bytes for <code>AF_INET6</code>.</p>



The `inet_addr()` and `inet_network()` routines interpret character strings representing numbers expressed in the IPv4 standard '.' notation, returning numbers suitable for use as IPv4 addresses and IPv4 network numbers, respectively. The routine `inet_makeaddr()` takes an IPv4 network number and a local network address and constructs an IPv4 address from it. The routines `inet_netof()` and `inet_lnaof()` break apart IPv4 host addresses, returning the network number and local network address part, respectively.

The `inet_ntoa()` routine returns a pointer to a string in the base 256 notation `d.d.d.d`. See INTERNET ADDRESSES.

Internet addresses are returned in network order, bytes ordered from left to right. Network numbers and local address parts are returned as machine format integer values.

## IPv6 Addresses

There are three conventional forms for representing IPv6 addresses as strings:

1. The preferred form is `x:x:x:x:x:x:x:x`, where the 'x's are the hexadecimal values of the eight 16-bit pieces of the address, for example,

```
1080:0:0:0:8:800:200C:417A
```

Note that it is not necessary to write the leading zeros in an individual field.

However, there must be at least one numeral in every field, except as described below.

2. Due to some methods of allocating certain styles of IPv6 addresses, it will be common for addresses to contain long strings of zero bits. In order to make writing addresses containing zero bits easier, a special syntax is available to compress the zeros. The use of ":" indicates multiple groups of 16-bits of zeros. The "::" can only appear once in an address. The "::" can also be used to compress the leading and/or trailing zeros in an address. For example,

```
1080::8:800:200C:417A
```

3. An alternative form that is sometimes more convenient when dealing with a mixed environment of IPv4 and IPv6 nodes is `x:x:x:x:x:x.d.d.d.d`, where the 'x's are the hexadecimal values of the six high-order 16-bit pieces of the address, and the 'd's are the decimal values of the four low-order 8-bit pieces of the standard IPv4 representation address, for example,

```
::FFFF:129.144.52.38
```

```
::129.144.52.38
```

where "`::FFFF:d.d.d.d`" and "`::d.d.d.d`" are, respectively, the general forms of an IPv4-mapped IPv6 address and an IPv4-compatible IPv6 address. Note that the IPv4 portion must be in the "`d.d.d.d`" form. The following forms are invalid:

```
::FFFF:d.d.d
```

```
::FFFF:d.d
```

```
::d.d.d
```

```
::d.d
```

The following form:

inet(3SOCKET)

`::FFFF:d`  
is valid, however it is an unconventional representation of the IPv4-compatible IPv6 address,  
`::255.255.0.d`  
while `:::d` corresponds to the general IPv6 address `"0:0:0:0:0:0:0:d"`.

#### IPv4 Addresses

Values specified using `'.'` notation take one of the following forms:

`d.d.d.d`  
`d.d.d`  
`d.d`

`d`When four parts are specified, each is interpreted as a byte of data and assigned, from left to right, to the four bytes of an IPv4 address.

When a three part address is specified, the last part is interpreted as a 16-bit quantity and placed in the right most two bytes of the network address. This makes the three part address format convenient for specifying Class B network addresses as `128.net.host`.

When a two part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the right most three bytes of the network address. This makes the two part address format convenient for specifying Class A network addresses as `net.host`.

When only one part is given, the value is stored directly in the network address without any byte rearrangement.

With the exception of `inet_pton()`, numbers supplied as *parts* in `'.'` notation may be decimal, octal, or hexadecimal, as specified in the C language. For example, a leading `0x` or `0X` implies hexadecimal; otherwise, a leading `0` implies octal; otherwise, the number is interpreted as decimal.

For IPv4 addresses, `inet_pton()` only accepts a string in the standard IPv4 dotted-decimal form:

`d.d.d.d`where each number has one to three digits with a decimal value between 0 and 255.

#### RETURN VALUES

The `inet_ntop()` routine returns a pointer to the buffer containing a string if the conversion succeeds, and `NULL` otherwise. Upon failure, `errno` is set to `EAFNOSUPPORT` if the *af* argument is invalid or `ENOSPC` if the size of the result buffer is inadequate.

`inet_pton()` returns 1 if the conversion succeeds, 0 if the input is not a valid IPv4 dotted-decimal string or a valid IPv6 address string, or -1 with `errno` set to `EAFNOSUPPORT` if the *af* argument is unknown.

The value -1 is returned by `inet_addr()` and `inet_network()` for malformed requests.

inet(3SOCKET)

The routines `inet_netof()` and `inet_lnaof()` break apart IPv4 host addresses, returning the network number and local network address part, respectively.

The routine `inet_ntoa()` returns a pointer to a string in the base 256 notation `d.d.d.d` described in INTERNET ADDRESSES.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `gethostbyname(3NSL)`, `getipnodebyname(3SOCKET)`, `getnetbyname(3SOCKET)`, `inet(3HEAD)`, `hosts(4)`, `ipnodes(4)`, `networks(4)`, `attributes(5)`

**NOTES** The return value from `inet_ntoa()` points to a buffer which is overwritten on each call. This buffer is implemented as thread-specific data in multithreaded applications.

**BUGS** The problem of host byte ordering versus network byte ordering is confusing. A simple way to specify Class C network addresses in a manner similar to that for Class B and Class A is needed.

## inet\_addr(3XNET)

<b>NAME</b>	inet_addr, inet_network, inet_makeaddr, inet_lnaof, inet_netof, inet_ntoa – Internet address manipulation						
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;arpa/inet.h&gt;  in_addr_t <b>inet_addr</b>(const char *cp); in_addr_t <b>inet_lnaof</b>(struct in_addr in); struct in_addr <b>inet_makeaddr</b>(in_addr_t net, in_addr_t lna); in_addr_t <b>inet_netof</b>(struct in_addr in); in_addr_t <b>inet_network</b>(const char *cp); char *<b>inet_ntoa</b>(struct in_addr in);</pre>						
<b>DESCRIPTION</b>	<p>The <code>inet_addr()</code> function converts the string pointed to by <i>cp</i>, in the Internet standard dot notation, to an integer value suitable for use as an Internet address.</p> <p>The <code>inet_lnaof()</code> function takes an Internet host address specified by <i>in</i> and extracts the local network address part, in host byte order.</p> <p>The <code>inet_makeaddr()</code> function takes the Internet network number specified by <i>net</i> and the local network address specified by <i>lna</i>, both in host byte order, and constructs an Internet address from them.</p> <p>The <code>inet_netof()</code> function takes an Internet host address specified by <i>in</i> and extracts the network number part, in host byte order.</p> <p>The <code>inet_network()</code> function converts the string pointed to by <i>cp</i>, in the Internet standard dot notation, to an integer value suitable for use as an Internet network number.</p> <p>The <code>inet_ntoa()</code> function converts the Internet host address specified by <i>in</i> to a string in the Internet standard dot notation.</p> <p>All Internet addresses are returned in network order (bytes ordered from left to right).</p> <p>Values specified using dot notation take one of the following forms:</p> <table><tr><td>a . b . c . d</td><td>When four parts are specified, each is interpreted as a byte of data and assigned, from left to right, to the four bytes of an Internet address.</td></tr><tr><td>a . b . c</td><td>When a three-part address is specified, the last part is interpreted as a 16-bit quantity and placed in the rightmost two bytes of the network address. This makes the three-part address format convenient for specifying Class B network addresses as <i>128 . net . host</i>.</td></tr><tr><td>a . b</td><td>When a two-part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the rightmost three bytes of the</td></tr></table>	a . b . c . d	When four parts are specified, each is interpreted as a byte of data and assigned, from left to right, to the four bytes of an Internet address.	a . b . c	When a three-part address is specified, the last part is interpreted as a 16-bit quantity and placed in the rightmost two bytes of the network address. This makes the three-part address format convenient for specifying Class B network addresses as <i>128 . net . host</i> .	a . b	When a two-part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the rightmost three bytes of the
a . b . c . d	When four parts are specified, each is interpreted as a byte of data and assigned, from left to right, to the four bytes of an Internet address.						
a . b . c	When a three-part address is specified, the last part is interpreted as a 16-bit quantity and placed in the rightmost two bytes of the network address. This makes the three-part address format convenient for specifying Class B network addresses as <i>128 . net . host</i> .						
a . b	When a two-part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the rightmost three bytes of the						

inet\_addr(3XNET)

network address. This makes the two-part address format convenient for specifying Class A network addresses as *net.host*.

a When only one part is given, the value is stored directly in the network address without any byte rearrangement.

All numbers supplied as parts in dot notation may be decimal, octal, or hexadecimal, that is, a leading 0x or 0X implies hexadecimal, as specified in the *ISO C* standard; otherwise, a leading 0 implies octal; otherwise, the number is interpreted as decimal.

**USAGE** The return value of `inet_ntoa()` may point to static data that may be overwritten by subsequent calls to `inet_ntoa()`.

**RETURN VALUES** Upon successful completion, `inet_addr()` returns the Internet address. Otherwise, it returns `(in_addr_t)(-1)`.

Upon successful completion, `inet_network()` returns the converted Internet network number. Otherwise, it returns `(in_addr_t)(-1)`.

The `inet_makeaddr()` function returns the constructed Internet address.

The `inet_lnaof()` function returns the local network address part.

The `inet_netof()` function returns the network number.

The `inet_ntoa()` function returns a pointer to the network address in Internet-standard dot notation.

**ERRORS** No errors are defined.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `endhostent(3XNET)`, `endnetent(3XNET)`, `attributes(5)`

## ldap(3LDAP)

<b>NAME</b>	ldap – Lightweight Directory Access Protocol package
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ] #include &lt;lber.h&gt; #include &lt;ldap.h&gt;</pre>
<b>DESCRIPTION</b>	<p>The Lightweight Directory Access Protocol (“LDAP”) package (SUNWlldap) includes various command line LDAP clients and a LDAP client library to provide programmatic access to the LDAP protocol. This man page gives an overview of the LDAP client library functions.</p> <p>An application might use the LDAP client library functions as follows. The application would initialize a LDAP session with a LDAP server by calling <code>ldap_init(3LDAP)</code>. Next, it authenticates to the LDAP server by calling <code>ldap_sasl_bind(3LDAP)</code> and friends. It may perform some LDAP operations and obtain results by calling <code>ldap_search(3LDAP)</code> and friends. To parse the results returned from these functions, it calls <code>ldap_parse_result(3LDAP)</code>, <code>ldap_next_entry(3LDAP)</code>, and <code>ldap_first_entry(3LDAP)</code> and others. It closes the LDAP session by calling <code>ldap_unbind(3LDAP)</code>.</p> <p>LDAP operations can be either synchronous or asynchronous. By convention, the names of the synchronous functions end with “_s.” For example, a synchronous binding to the LDAP server can be performed by calling <code>tlldap_sasl_bind_s(3LDAP)</code>. Complete an asynchronous binding with <code>ldap_sasl_bind(3LDAP)</code>. All synchronous functions return the actual outcome of the operation, either <code>LDAP_SUCCESS</code> or an error code. Asynchronous routines provide an invocation identifier which can be used to obtain the result of a specific operation by passing it to the <code>tlldap_result(3LDAP)</code> function.</p>
<b>Initializing a LDAP session</b>	<p>Initializing a LDAP session involves calling the <code>ldap_init(3LDAP)</code> function. However, the call does not actually open a connection to the LDAP server. It merely initializes a LDAP structure that represents the session. The connection is opened when the first operation is attempted. Unlike <code>ldap_init()</code>, <code>ldap_open(3LDAP)</code> attempts to open a connection with the LDAP server. However, the use of <code>ldap_open()</code> is deprecated.</p>
<b>Authenticating to a LDAP server</b>	<p>The <code>ldap_sasl_bind(3LDAP)</code> and <code>ldap_sasl_bind_s(3LDAP)</code> functions provide general and extensible authentication for an LDAP client to a LDAP server. Both use the Simple Authentication Security Layer (SASL). Simplified routines <code>ldap_simple_bind(3LDAP)</code> and <code>ldap_simple_bind_s(3LDAP)</code> use cleartext passwords to bind to the LDAP server. Use of <code>ldap_bind(3LDAP)</code> and <code>ldap_bind_s(3LDAP)</code> is deprecated.</p>
<b>Searching a LDAP directory</b>	<p>Search for an entry in a LDAP directory by calling the <code>ldap_search_ext(3LDAP)</code> or the <code>ldap_search_ext_s(3LDAP)</code> functions. These functions support LDAPv3 server controls, client controls and variable size and time limits as arguments for each search operation. <code>ldap_search(3LDAP)</code> and <code>ldap_search_s(3LDAP)</code> are identical functions but do not support the controls and limits as arguments to the call.</p>

<b>Adding or Deleting an entry</b>	Use <code>ldap_add_ext(3LDAP)</code> and <code>ldap_delete_ext(3LDAP)</code> to add or delete entries in a LDAP directory server. The synchronous counterparts to these functions are <code>ldap_add_ext_s(3LDAP)</code> and <code>ldap_delete_ext_s(3LDAP)</code> . The <code>ldap_add(3LDAP)</code> , <code>ldap_add_s(3LDAP)</code> , <code>ldap_delete(3LDAP)</code> , and <code>ldap_delete_s(3LDAP)</code> provide identical functionality to add and to delete entries, but they do not support LDAP v3 server and client controls.
<b>Modifying Entries</b>	Use <code>ldap_modify_ext(3LDAP)</code> and <code>ldap_modify_ext_s(3LDAP)</code> to modify an existing entry in a LDAP server that supports for LDAPv3 server and client controls. Similarly, use <code>ldap_rename(3LDAP)</code> and <code>ldap_rename_s(3LDAP)</code> to change the name of an LDAP entry. The <code>ldap_modrdn(3LDAP)</code> , <code>ldap_modrdn_s()</code> , <code>ldap_modrdn2(3LDAP)</code> and <code>ldap_modrdn2_s(3LDAP)</code> interfaces are deprecated.
<b>Obtaining Results</b>	Use <code>ldap_result(3LDAP)</code> to obtain the results of a previous asynchronous operation. For all LDAP operations other than search, only one message is returned. For the search operation, a list of result messages can be returned.
<b>Handling Errors and Parsing Results</b>	Use the <code>ldap_parse_result(3LDAP)</code> , <code>ldap_parse_sasl_bind_result(3LDAP)</code> , and the <code>ldap_parse_extended_result(3LDAP)</code> functions to extract required information from results and to handle the returned errors. To convert a numeric error code into a null-terminated character string message describing the error, use <code>ldap_err2string(3LDAP)</code> . The <code>ldap_result2error(3LDAP)</code> and <code>ldap_perror(3LDAP)</code> functions are deprecated. To step through the list of messages in a result returned by <code>ldap_result()</code> , use <code>ldap_first_message(3LDAP)</code> and <code>ldap_next_message(3LDAP)</code> . <code>ldap_count_messages(3LDAP)</code> returns the number of messages contained in the list.  You can use <code>ldap_first_entry(3LDAP)</code> and <code>ldap_next_entry(3LDAP)</code> to step through and obtain a list of entries from a list of messages returned by a search result. <code>ldap_count_entries(3LDAP)</code> returns the number of entries contained in a list of messages. Call either <code>ldap_first_attribute(3LDAP)</code> and <code>ldap_next_attribute(3LDAP)</code> to step through a list of attributes associated with an entry. Retrieve the values of a given attribute by calling <code>ldap_get_values(3LDAP)</code> and <code>ldap_get_values_len(3LDAP)</code> . Count the number of values returned by using <code>ldap_count_values(3LDAP)</code> and <code>ldap_count_values_len(3LDAP)</code> .  Use the <code>ldap_get_lang_values(3LDAP)</code> and <code>ldap_get_lang_values_len(3LDAP)</code> to return an attribute's values that matches a specified language subtype. The <code>ldap_get_lang_values()</code> function returns an array of an attribute's string values that matches a specified language subtype. To retrieve the binary data from an attribute, call the <code>ldap_get_lang_values_len()</code> function instead.
<b>Uniform Resource Locators (URLS)</b>	You can use the <code>ldap_url(3LDAP)</code> functions to test a URL to verify that it is an LDAP URL, to parse LDAP URLs into their component pieces, to initiate searches directly using an LDAP URL, and to retrieve the URL associated with a DNS domain name or a distinguished name.

## ldap(3LDAP)

<b>User Friendly Naming</b>	The <code>ldap_ufn(3LDAP)</code> functions implement a user friendly naming scheme by means of LDAP. This scheme allows you to look up entries using fuzzy, untyped names like "mark smith, umich, us".
<b>Caching</b>	The <code>ldap_memcache(3LDAP)</code> functions provide an in-memory client side cache to store search requests. Caching improves performance and reduces network bandwidth when a client makes repeated requests.
<b>Utility Functions</b>	There are also various utility functions. You can use the <code>ldap_sort(3LDAP)</code> functions are used to sort the entries and values returned by means of the ldap search functions. The <code>ldap_friendly(3LDAP)</code> functions will map from short two letter country codes or other strings to longer "friendlier" names. Use the <code>ldap_charset(3LDAP)</code> functions to translate to and from the T.61 character set that is used for many character strings in the LDAP protocol.
<b>Generating Filters</b>	Make calls to <code>ldap_init_getfilter(3LDAP)</code> and <code>ldap_init_getfilter_buf()</code> to generate filters to be used in <code>ldap_search(3LDAP)</code> and <code>ldap_search_s(3LDAP)</code> . <code>ldap_init_getfilter()</code> reads <code>ldapfilter.conf(4)</code> , the LDAP configuration file, while <code>ldap_init_getfilter_buf()</code> reads the configuration information from <i>buf</i> of length <i>buflen</i> . <code>ldap_getfilter_free(3LDAP)</code> frees memory that has been allocated by means of <code>ldap_init_getfilter()</code> .
<b>BER Library</b>	The LDAP package includes a set of lightweight Basic Encoding Rules ("BER") functions. The LDAP library functions use the BER functions to encode and decode LDAP protocol elements through the slightly simplified BER defined by LDAP. They are not normally used directly by An LDAP application program will not normally use the BER functions directly. Instead, these functions provide a <code>printf()</code> and <code>scanf()</code> -like interface, as well as lower-level access.
<b>LIST OF INTERFACES</b>	<code>ldap_open(3LDAP)</code> Deprecated. Use <code>ldap_init(3LDAP)</code> . <code>ldap_init(3LDAP)</code> Initialize a session with a LDAP server without opening a connection to a server. <code>ldap_result(3LDAP)</code> Obtain the result from a previous asynchronous operation. <code>ldap_abandon(3LDAP)</code> Abandon or abort an asynchronous operation. <code>ldap_add(3LDAP)</code> Asynchronously add an entry <code>ldap_add_s(3LDAP)</code> Synchronously add an entry. <code>ldap_add_ext(3LDAP)</code> Asynchronously add an entry with support for LDAPv3 controls. <code>ldap_add_ext_s(3LDAP)</code> Synchronously add an entry with support for LDAPv3 controls.



`ldap_bind(3LDAP)`  
 Deprecated. Use `ldap_sasl_bind(3LDAP)` or `ldap_simple_bind(3LDAP)`.

`ldap_sasl_bind(3LDAP)`  
 Asynchronously bind to the directory using SASL authentication

`ldap_sasl_bind_s(3LDAP)`  
 Synchronously bind to the directory using SASL authentication

`ldap_bind_s(3LDAP)`  
 Deprecated. Use `ldap_sasl_bind_s(3LDAP)` or `ldap_simple_bind_s(3LDAP)`.

`ldap_simple_bind(3LDAP)`  
 Asynchronously bind to the directory using simple authentication.

`ldap_simple_bind_s(3LDAP)`  
 Synchronously bind to the directory using simple authentication.

`ldap_unbind(3LDAP)`  
 Synchronously unbind from the LDAP server, close the connection, and dispose the session handle.

`ldap_unbind_ext(3LDAP)`  
 Synchronously unbind from the LDAP server and close the connection.  
`ldap_unbind_ext()` allows you to explicitly include both server and client controls in the unbind request.

`ldap_set_rebind_proc(3LDAP)`  
 Set callback function for obtaining credentials from a referral.

`ldap_memcache_init(3LDAP)`  
 Create the in-memory client side cache.

`ldap_memcache_set(3LDAP)`  
 Associate an in-memory cache that has been already created by calling the `ldap_memcache_init(3LDAP)` function with an LDAP connection handle.

`ldap_memcache_get(3LDAP)`  
 Get the cache associated with the specified LDAP structure.

`ldap_memcache_flush(3LDAP)`  
 Flushes search requests from the cache.

`ldap_memcache_destroy(3LDAP)`  
 Frees the specified `LDAPMemCache` structure pointed to by cache from memory.

`ldap_memcache_update(3LDAP)`  
 Checks the cache for items that have expired and removes them.

`ldap_compare(3LDAP)`  
 Asynchronous compare with a directory entry.

`ldap_compare_s(3LDAP)`  
 Synchronous compare with a directory entry.

## ldap(3LDAP)

`ldap_compare_ext(3LDAP)`  
Asynchronous compare with a directory entry, with support for LDAPv3 controls.

`ldap_compare_ext_s(3LDAP)`  
Synchronous compare with a directory entry, with support for LDAPv3 controls.

`ldap_control_free(3LDAP)`  
Dispose of an LDAP control.

`ldap_controls_free(3LDAP)`  
Dispose of an array of LDAP controls.

`ldap_delete(3LDAP)`  
Asynchronously delete an entry.

`ldap_delete_s(3LDAP)`  
Synchronously delete an entry.

`ldap_delete_ext(3LDAP)`  
Asynchronously delete an entry, with support for LDAPv3 controls.

`ldap_delete_ext_s(3LDAP)`  
Synchronously delete an entry, with support for LDAPv3 controls.

`ldap_init_templates(3LDAP)`  
Read a sequence of templates from a LDAP template configuration file.

`ldap_init_templates_buf(3LDAP)`  
Read a sequence of templates from a buffer.

`ldap_free_templates(3LDAP)`  
Dispose of the templates allocated.

`ldap_first_reference(3LDAP)`  
Step through a list of continuation references from a search result.

`ldap_next_reference(3LDAP)`  
Step through a list of continuation references from a search result.

`ldap_count_references(3LDAP)`  
Count the number of messages in a search result.

`ldap_first_message(3LDAP)`  
Step through a list of messages in a search result.

`ldap_count_messages(3LDAP)`  
Count the messages in a list of messages in a search result.

`ldap_next_message(3LDAP)`  
Step through a list of messages in a search result.

`ldap_msgtype(3LDAP)`  
Return the type of LDAP message.

`ldap_first_disptmpl(3LDAP)`  
Get first display template in a list.

`ldap_next_disptmpl(3LDAP)`  
Get next display template in a list.

`ldap_oc2template(3LDAP)`  
Return template appropriate for the objectclass.

`ldap_name2template(3LDAP)`  
Return named template

`ldap_tmplatattrs(3LDAP)`  
Return attributes needed by the template.

`ldap_first_tmplrow(3LDAP)`  
Return first row of displayable items in a template.

`ldap_next_tmplrow(3LDAP)`  
Return next row of displayable items in a template.

`ldap_first_tmplcol(3LDAP)`  
Return first column of displayable items in a template.

`ldap_next_tmplcol(3LDAP)`  
Return next column of displayable items in a template.

`ldap_entry2text(3LDAP)`  
Display an entry as text by using a display template.

`ldap_entry2text_search(3LDAP)`  
Search for and display an entry as text by using a display template.

`ldap_vals2text(3LDAP)`  
Display values as text.

`ldap_entry2html(3LDAP)`  
Display an entry as HTML (HyperText Markup Language) by using a display template.

`ldap_entry2html_search(3LDAP)`  
Search for and display an entry as HTML by using a display template.

`ldap_vals2html(3LDAP)`  
Display values as HTML.

`ldap_perror(3LDAP)`  
Deprecated. Use `ldap_parse_result(3LDAP)`.

`ldap_result2error(3LDAP)`  
Deprecated. Use `ldap_parse_result(3LDAP)`.

`ldap_err2string(3LDAP)`  
Convert LDAP error indication to a string.

`ldap_first_attribute(3LDAP)`  
Return first attribute name in an entry.

## ldap(3LDAP)

`ldap_next_attribute(3LDAP)`  
Return next attribute name in an entry.

`ldap_first_entry(3LDAP)`  
Return first entry in a chain of search results.

`ldap_next_entry(3LDAP)`  
Return next entry in a chain of search results.

`ldap_count_entries(3LDAP)`  
Return number of entries in a search result.

`ldap_friendly_name(3LDAP)`  
Map from unfriendly to friendly names.

`ldap_free_friendlymap(3LDAP)`  
Free resources used by `ldap_friendly(3LDAP)`.

`ldap_get_dn(3LDAP)`  
Extract the DN from an entry.

`ldap_explode_dn(3LDAP)`  
Convert a DN into its component parts.

`ldap_explode_dns(3LDAP)`  
Convert a DNS-style DN into its component parts (experimental).

`ldap_is_dns_dn(3LDAP)`  
Check to see if a DN is a DNS-style DN (experimental).

`ldap_dns_to_dn(3LDAP)`  
Convert a DNS domain name into an X.500 distinguished name.

`ldap_dn2ufn(3LDAP)`  
Convert a DN into user friendly form.

`ldap_get_values(3LDAP)`  
Return an attribute's values.

`ldap_get_values_len(3LDAP)`  
Return an attribute's values with lengths.

`ldap_value_free(3LDAP)`  
Free memory allocated by `ldap_get_values(3LDAP)`.

`ldap_value_free_len(3LDAP)`  
Free memory allocated by `ldap_get_values_len(3LDAP)`.

`ldap_count_values(3LDAP)`  
Return number of values.

`ldap_count_values_len(3LDAP)`  
Return number of values.

`ldap_init_getfilter(3LDAP)`  
Initialize getfilter functions from a file.

`ldap_init_getfilter_buf(3LDAP)`  
Initialize getfilter functions from a buffer.

`ldap_getfilter_free(3LDAP)`  
Free resources allocated by `ldap_init_getfilter(3LDAP)`.

`ldap_getfirstfilter(3LDAP)`  
Return first search filter.

`ldap_getnextfilter(3LDAP)`  
Return next search filter.

`ldap_build_filter(3LDAP)`  
Construct an LDAP search filter from a pattern.

`ldap_setfilteraffixes(3LDAP)`  
Set prefix and suffix for search filters.

`ldap_modify(3LDAP)`  
Asynchronously modify an entry.

`ldap_modify_s(3LDAP)`  
Synchronously modify an entry.

`ldap_modify_ext(3LDAP)`  
Asynchronously modify an entry, return value, and place message.

`ldap_modify_ext_s(3LDAP)`  
Synchronously modify an entry, return value, and place message.

`ldap_mods_free(3LDAP)`  
Free array of pointers to mod structures used by `ldap_modify(3LDAP)`.

`ldap_modrdn2(3LDAP)`  
Deprecated. Use `ldap_rename(3LDAP)` instead.

`ldap_modrdn2_s(3LDAP)`  
Deprecated. Use `ldap_rename_s(3LDAP)` instead.

`ldap_modrdn(3LDAP)`  
Deprecated. Use `ldap_rename(3LDAP)` instead.

`ldap_modrdn_s(3LDAP)`  
Deprecated. Use `ldap_rename_s(3LDAP)` instead.

`ldap_rename(3LDAP)`  
Asynchronously modify the name of an LDAP entry.

`ldap_rename_s(3LDAP)`  
Synchronously modify the name of an LDAP entry.

`ldap_msgfree(3LDAP)`  
Free result messages.

`ldap_parse_result(3LDAP)`  
Search for a message to parse.

## ldap(3LDAP)

`ldap_parse_extended_result(3LDAP)`  
Search for a message to parse.

`ldap_parse_sasl_bind_result(3LDAP)`  
Search for a message to parse.

`ldap_search(3LDAP)`  
Asynchronously search the directory.

`ldap_search_s(3LDAP)`  
Synchronously search the directory.

`ldap_search_ext(3LDAP)`  
Asynchronously search the directory with support for LDAPv3 controls.

`ldap_search_ext_s(3LDAP)`  
Synchronously search the directory with support for LDAPv3 controls.

`ldap_search_st(3LDAP)`  
Synchronously search the directory with support for a local timeout value.

`ldap_ufn_search_s(3LDAP)`  
User friendly search the directory.

`ldap_ufn_search_c(3LDAP)`  
User friendly search the directory with cancel.

`ldap_ufn_search_ct(3LDAP)`  
User friendly search the directory with cancel and timeout.

`ldap_ufn_setfilter(3LDAP)`  
Set filter file used by `ldap_ufn(3LDAP)` functions.

`ldap_ufn_setprefix(3LDAP)`  
Set prefix used by `ldap_ufn(3LDAP)` functions.

`ldap_ufn_timeout(3LDAP)`  
Set timeout used by `ldap_ufn(3LDAP)` functions.

`ldap_is_ldap_url(3LDAP)`  
Check a URL string to see if it is an LDAP URL.

`ldap_url_parse(3LDAP)`  
Break up an LDAP URL string into its components.

`ldap_free_urldesc(3LDAP)`  
Free an LDAP URL structure.

`ldap_url_search(3LDAP)`  
Asynchronously search by using an LDAP URL.

`ldap_url_search_s(3LDAP)`  
Synchronously search by using an LDAP URL.

`ldap_url_search_st(3LDAP)`  
Asynchronously search by using an LDAP URL, with support for a local timeout value.

`ldap_dns_to_url(3LDAP)`  
Locate the LDAP URL associated with a DNS domain name.

`ldap_dn_to_url(3LDAP)`  
Locate the LDAP URL associated with a distinguished name.

`ldap_init_searchprefs(3LDAP)`  
Initialize searchprefs functions from a file.

`ldap_init_searchprefs_buf(3LDAP)`  
Initialize searchprefs functions from a buffer.

`ldap_free_searchprefs(3LDAP)`  
Free memory allocated by searchprefs functions.

`ldap_first_searchobj(3LDAP)`  
Return first searchpref object.

`ldap_next_searchobj(3LDAP)`  
Return next searchpref object.

`ldap_sort_entries(3LDAP)`  
Sort a list of search results.

`ldap_sort_values(3LDAP)`  
Sort a list of attribute values.

`ldap_sort_strcasecmp(3LDAP)`  
Case insensitive string comparison.

`ldap_set_string_translators(3LDAP)`  
Set character set translation functions used by LDAP library.

`ldap_translate_from_t61(3LDAP)`  
Translate from the T.61 character set to another character set.

`ldap_translate_to_t61(3LDAP)`  
Translate to the T.61 character set from another character set.

`ldap_enable_translation(3LDAP)`  
Enable or disable character translation for an LDAP entry result.

`ldap_version(3LDAP)`  
Get version information about the LDAP SDK for C.

`ldap_get_lang_values(3LDAP)`  
Return an attribute's value that matches a specified language subtype.

`ldap_get_lang_values_len(3LDAP)`  
Return an attribute's value that matches a specified language subtype along with lengths.

## ldap(3LDAP)

`ldap_get_entry_controls(3LDAP)`

Get the LDAP controls included with a directory entry in a set of search results.

`ldap_get_option(3LDAP)`

Get session preferences in an LDAP structure.

`ldap_set_option(3LDAP)`

Set session preferences in an LDAP structure.

`ldap_memfree(3LDAP)`

Free memory allocated by LDAP API functions.

### ATTRIBUTES

See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Stability Level	Evolving

### SEE ALSO

`attributes(5)`



**NAME** ldap\_abandon – abandon an LDAP operation in progress

**SYNOPSIS**

```
cc [ flag... ] file... -lldap [ library... ]

#include <lber.h>
#include <ldap.h>

int ldap_abandon(LDAP *ld, int msgid);
```

**DESCRIPTION** The ldap\_abandon() function is used to abandon or cancel an LDAP operation in progress. The *msgid* passed should be the message id of an outstanding LDAP operation, as returned by ldap\_search(3LDAP), ldap\_modify(3LDAP), etc.

ldap\_abandon() checks to see if the result of the operation has already come in. If it has, it deletes it from the queue of pending messages. If not, it sends an LDAP abandon operation to the the LDAP server.

The caller can expect that the result of an abandoned operation will not be returned from a future call to ldap\_result(3LDAP).

**ERRORS** ldap\_abandon() returns 0 if successful or -1 otherwise and setting *ld\_errno* appropriately. See ldap\_error(3LDAP) for details.

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_result(3LDAP), ldap\_error(3LDAP), attributes(5)

## ldap\_add(3LDAP)

<b>NAME</b>	ldap_add, ldap_add_s, ldap_add_ext, ldap_add_ext_s – perform an LDAP add operation
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_add(LDAP *ld, char *dn, LDAPMod *attrs[]); int ldap_add_s(LDAP *ld, char *dn, LDAPMod *attrs[]); int ldap_add_ext(LDAP *ld, char *dn, LDAPMod **attrs, LDAPControl **serverctrls, int *msgidp); int ldap_add_ext_s(LDAP *ld, char *dn, LDAPMod **attrs, LDAPControl **serverctrls, LDAPControl **clientctrls);</pre>
<b>DESCRIPTION</b>	<p>The <code>ldap_add_s()</code> function is used to perform an LDAP add operation. It takes <i>dn</i>, the DN of the entry to add, and <i>attrs</i>, a null-terminated array of the entry's attributes. The LDAPMod structure is used to represent attributes, with the <i>mod_type</i> and <i>mod_values</i> fields being used as described under <code>ldap_modify(3LDAP)</code>, and the <i>ldap_op</i> field being used only if you need to specify the LDAP_MOD_BVALUES option. Otherwise, it should be set to zero.</p> <p>Note that all entries except that specified by the last component in the given DN must already exist. <code>ldap_add_s()</code> returns an LDAP error code indicating success or failure of the operation. See <code>ldap_error(3LDAP)</code> for more details.</p> <p>The <code>ldap_add()</code> function works just like <code>ldap_add_s()</code>, but it is asynchronous. It returns the message id of the request it initiated. The result of this operation can be obtained by calling <code>ldap_result(3LDAP)</code>.</p> <p>The <code>ldap_add_ext()</code> function initiates an asynchronous add operation and returns LDAP_SUCCESS if the request was successfully sent to the server, or else it returns a LDAP error code if not (see <code>ldap_error(3LDAP)</code>). If successful, <code>ldap_add_ext()</code> places the message id of <i>msgidp</i>. A subsequent call to <code>ldap_result()</code>, can be used to obtain the result of the add request.</p> <p>The <code>ldap_add_ext_s()</code> function initiates a synchronous add operation and returns the result of the operation itself.</p>
<b>ERRORS</b>	<code>ldap_add()</code> returns -1 in case of error initiating the request, and will set the <i>ld_errno</i> field in the <i>ld</i> parameter to indicate the error. <code>ldap_add_s()</code> will return an LDAP error code directly.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for a description of the following attributes:

ldap\_add(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_error(3LDAP), ldap\_modify(3LDAP), attributes(5)

## ldap\_ber\_free(3LDAP)

**NAME** ldap\_ber\_free – free a BerElement structure from memory

**SYNOPSIS** cc -flag ... file ...-lldap [-library ...]  
#include <ldap.h>

```
void ldap_ber_free(BerElement *ber, int freebuf);
```

**DESCRIPTION** You can make a call to the ldap\_ber\_free() function to free BerElement structures allocated by ldap\_first\_attribute() and by ldap\_next\_attribute() function calls. When freeing structures allocated by these functions, specify 0 for the freebuf argument. The ldap\_first\_attribute() and by ldap\_next\_attribute() functions do not allocate the extra buffer in the BerElement structure.

For example, to retrieve attributes from a search result entry, you need to call the ldap\_first\_attribute() function. A call to this function allocates a BerElement structure, which is used to help track the current attribute. When you are done working with the attributes, this structure should be freed from memory, if it still exists.

This function is deprecated . Use the ber\_free() function instead.

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Obsolete

**SEE ALSO** ber\_free(3LDAP), ldap\_first\_attribute(3LDAP), ldap\_next\_attribute(3LDAP), attributes(5)

<b>NAME</b>	ldap_bind, ldap_bind_s, ldap_sasl_bind, ldap_sasl_bind_s, ldap_simple_bind, ldap_simple_bind_s, ldap_unbind, ldap_unbind_s, ldap_unbind_ext, ldap_set_rebind_proc – LDAP bind functions
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap[ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_bind(LDAP *ld, char *who, char *cred, int method); int ldap_bind_s(LDAP *ld, char *who, char *cred, int method); int ldap_simple_bind(LDAP *ld, char *who, char *passwd); int ldap_simple_bind_s(LDAP *ld, char *who, char *passwd); int ldap_unbind(LDAP *ld); int ldap_unbind_s(LDAP *ld); int ldap_unbind_ext(LDAP *ld, LDAPControl **serverctrls, LDAPControl **clientctrls);  void ldap_set_rebind_proc(LDAP *ld, int (*rebindproc));  int ldap_sasl_bind(LDAP *ld, char *dn, char *mechanism, struct berval **serverctrls, LDAPControl **clientctrls, int *msgidp);  int ldap_sasl_bind_s(LDAP *ld, char *dn, char *mechanism, struct berval *cred, LDAPControl **serverctrls, LDAPControl **clientctrls);</pre>
<b>DESCRIPTION</b>	<p>These functions provide various interfaces to the LDAP bind operation. After a connection is made to an LDAP server using <code>ldap_open(3LDAP)</code>, an LDAP bind operation must be performed before other operations can be attempted over the connection. Both synchronous and asynchronous versions of each variant of the bind call are provided. There are three types of bind calls, simple authentication, kerberos authentication, and general functions. All LDAP bind functions take <code>ld</code> as their first parameter, which is returned from <code>ldap_open(3LDAP)</code>.</p>
<b>Simple Authentication</b>	<p>The simplest form of the bind call is <code>ldap_simple_bind_s()</code>. It takes the DN to bind as in <code>who</code>, and the user password associated with the entry in <code>passwd</code>. It returns an LDAP error code. See <code>ldap_error(3LDAP)</code>. The <code>ldap_simple_bind()</code> call is asynchronous. It takes the same parameters but only initiates the bind operation and returns the message id of the request it sent. You can obtain the result of the operation by a subsequent call to <code>ldap_result(3LDAP)</code>.</p>
<b>General Authentication</b>	<p>Use the <code>ldap_bind()</code> and <code>ldap_bind_s()</code> functions when the authentication method to use needs to be selected at runtime. Both functions take an extra <code>method</code> parameter that selects the authentication method to use. It should be set to <code>LDAP_AUTH_SIMPLE</code> to select simple authentication. <code>ldap_bind()</code> returns the message id of the request it initiates. <code>ldap_bind_s()</code> returns an LDAP error code.</p>

## ldap\_bind(3LDAP)

	<p>The <code>ldap_sasl_bind()</code> and <code>ldap_sasl_bind_s()</code> functions are used for general and extensible authentication over LDAP through the use of the Simple Authentication Security Layer. The routines both take the dn to bind as the method to use. A dotted-string representation of an OID identifies the method, and a struct <code>berval</code> holds the credentials. The special constant value <code>LDAP_SASL_SIMPLE("")</code> can be passed to request simple authentication, or the simplified routines <code>ldap_simple_bind()</code> or <code>ldap_simple_bind_s()</code> can be use.</p>
<b>Unbinding</b>	<p>The <code>ldap_unbind()</code> call is used to unbind from the directory, terminate the current association, and free the resources contained in the <code>ld</code> structure. Once it is called, the connection to the LDAP server is closed, and the <code>ld</code> structure is invalid. The <code>ldap_unbind_s()</code> call is just another name for <code>ldap_unbind()</code>. Both of these calls are synchronous in nature.</p> <p>The <code>ldap_unbind_ext()</code> function unbinds from the directory, terminates the current association, and frees the resources contained in the LDAP structure. Unlike <code>ldap_unbind()</code> and <code>ldap_unbind_s()</code>, you can explicitly include both server and client controls with the <code>with_ldap_unbind_ext()</code> request. Since there is no server response to an unbind request, you will not receive a response from a server control that is included with the unbind request.</p>
<b>Rebinding While Following Referral</b>	<p>The <code>ldap_set_rebind_proc()</code> call is used to set a function that will be called back to obtain bind credentials used when a new server is contacted following an LDAP referral. If <code>ldap_set_rebind_proc()</code> is never called, or if it is called with a NULL <code>rebindproc</code> parameter, an unauthenticated simple LDAP bind will always be done when chasing referrals.</p> <p><code>rebindproc()</code> should be a function that is declared like this:</p> <pre>int rebindproc(LDAP *ld, char **whop, char **credp,                int *methodp, int freeit);</pre> <p>The LDAP library will first call the <code>rebindproc()</code> to obtain the referral bind credentials, and the <code>freeit</code> parameter will be zero. The <code>whop</code>, <code>credp</code>, and <code>methodp</code> should be set as appropriate. If <code>rebindproc()</code> returns <code>LDAP_SUCCESS</code>, referral processing continues, and the <code>rebindproc()</code> will be called a second time with <code>freeit</code> non-zero to give your application a chance to free any memory allocated in the previous call.</p> <p>If anything but <code>LDAP_SUCCESS</code> is returned by the first call to <code>rebindproc()</code>, then referral processing is stopped, and that error code is returned for the original LDAP operation.</p>
<b>RETURN VALUES</b>	Make a call to <code>ldap_result(3LDAP)</code> to obtain the result of a bind operation.
<b>ERRORS</b>	Asynchronous functions will return <code>-1</code> in case of error. See <code>ldap_error(3LDAP)</code> for more information on error codes returned.. If no credentials are returned, the result parameter is set to <code>NULL</code> .
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for a description of the following attributes:

ldap\_bind(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Stability Level	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_error(3LDAP), ldap\_open(3LDAP), attributes(5)

## ldap\_charset(3LDAP)

<b>NAME</b>	ldap_charset, ldap_set_string_translators, ldap_t61_to_8859, ldap_8859_to_t61, ldap_translate_from_t61, ldap_translate_to_t61, ldap_enable_translation – LDAP character set translation functions
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  void ldap_set_string_translators(LDAP *ld, BERTranslateProc     encode_proc, BERTranslateProc decode_proc);  typedef int (*BERTranslateProc)(char **bufp, unsigned long *buflenp,     int free_input);  int ldap_t61_to_8859(char **bufp, unsigned long *buflenp, int     free_input);  int ldap_8859_to_t61(char **bufp, unsigned long *buflenp, int     free_input);  int ldap_translate_from_t61(LDAP *ld, char **bufp, unsigned long     *lenp, int free_input);  int ldap_translate_to_t61(LDAP *ld, char **bufp, unsigned long *lenp,     int free_input);  void ldap_enable_translation(LDAP *ld, LDAPMessage *entry, int     enable);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to enable translation of character strings used in the LDAP library to and from the T.61 character set used in the LDAP protocol. These functions are only available if the LDAP and LBER libraries are compiled with STR_TRANSLATION defined. It is also possible to turn on character translation by default so that all LDAP library callers will experience translation; see the LDAP Make-common source file for details.</p> <p>ldap_set_string_translators() sets the translation functions that will be used by the LDAP library. They are not actually used until the <i>ld_lberoptions</i> field of the LDAP structure is set to include the LBER_TRANSLATE_STRINGS option.</p> <p>ldap_t61_to_8859() and ldap_8859_to_t61() are translation functions for converting between T.61 characters and ISO-8859 characters. The specific 8859 character set used is determined at compile time.</p> <p>ldap_translate_from_t61() is used to translate a string of characters from the T.61 character set to a different character set. The actual translation is done using the <i>decode_proc</i> that was passed to a previous call to ldap_set_string_translators(). On entry, <i>bufp</i> should point to the start of the T.61 characters to be translated and <i>lenp</i> should contain the number of bytes to translate. If <i>free_input</i> is non-zero, the input buffer will be freed if translation is a</p>



## ldap\_charset(3LDAP)

success. If the translation is a success, `LDAP_SUCCESS` will be returned, *\*bufp* will point to a newly malloc'd buffer that contains the translated characters, and *\*lenp* will contain the length of the result. If translation fails, an LDAP error code will be returned.

`ldap_translate_to_t61()` is used to translate a string of characters to the T.61 character set from a different character set. The actual translation is done using the *encode\_proc* that was passed to a previous call to `ldap_set_string_translators()`. This function is called just like `ldap_translate_from_t61()`.

`ldap_enable_translation()` is used to turn on or off string translation for the LDAP entry *entry* (typically obtained by calling `ldap_first_entry()` or `ldap_next_entry()` after a successful LDAP search operation). If *enable* is zero, translation is disabled; if non-zero, translation is enabled. This function is useful if you need to ensure that a particular attribute is not translated when it is extracted using `ldap_get_values()` or `ldap_get_values_len()`. For example, you would not want to translate a binary attributes such as `jpegPhoto`.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap(3LDAP)`, `attributes(5)`

## ldap\_compare(3LDAP)

<b>NAME</b>	ldap_compare, ldap_compare_s, ldap_compare_ext, ldap_compare_ext_s – LDAP compare operation				
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_compare(LDAP *ld, char *dn, char *attr, char *value);  int ldap_compare_s(LDAP *ld, char *dn, char *attr, char *value);  int ldap_compare_ext(LDAP *ld, char *dn, char *attr, struct berval     *bvalue, LDAPControl **serverctrls, LDAPControl **clientctrls, int     *msgidp);  int ldap_compare_ext_s(LDAP *ld, char *dn, char *attr, struct berval     *bvalue, LDAPControl **serverctrls, LDAPControl **clientctrls);</pre>				
<b>DESCRIPTION</b>	<p>The <code>ldap_compare_s()</code> function is used to perform an LDAP compare operation synchronously. It takes <code>dn</code>, the DN of the entry upon which to perform the compare, and <code>attr</code> and <code>value</code>, the attribute type and value to compare to those found in the entry. It returns an LDAP error code, which will be <code>LDAP_COMPARE_TRUE</code> if the entry contains the attribute value and <code>LDAP_COMPARE_FALSE</code> if it does not. Otherwise, some error code is returned.</p> <p>The <code>ldap_compare()</code> function is used to perform an LDAP compare operation asynchronously. It takes the same parameters as <code>ldap_compare_s()</code>, but returns the message id of the request it initiated. The result of the compare can be obtained by a subsequent call to <code>ldap_result(3LDAP)</code>.</p> <p>The <code>ldap_compare_ext()</code> function initiates an asynchronous compare operation and returns <code>LDAP_SUCCESS</code> if the request was successfully sent to the server, or else it returns a LDAP error code if not (see <code>ldap_error(3LDAP)</code>). If successful, <code>ldap_compare_ext()</code> places the message id of the request in <code>*msgidp</code>. A subsequent call to <code>ldap_result()</code>, can be used to obtain the result of the add request.</p> <p>The <code>ldap_compare_ext_s()</code> function initiates a synchronous compare operation and as such returns the result of the operation itself.</p>				
<b>ERRORS</b>	<code>ldap_compare_s()</code> returns an LDAP error code which can be interpreted by calling one of <code>ldap_perror(3LDAP)</code> and friends. <code>ldap_compare()</code> returns <code>-1</code> if something went wrong initiating the request. It returns the non-negative message id of the request if it was successful.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for a description of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWcsl (32-bit)</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsl (32-bit)
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsl (32-bit)				

ldap\_compare(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
	SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_error(3LDAP), attributes(5)

**BUGS** There is no way to compare binary values using ldap\_compare().

## ldap\_control\_free(3LDAP)

**NAME** ldap\_control\_free, ldap\_controls\_free – LDAP control disposal

**SYNOPSIS**

```
cc [ flag... ] file... -lldap [ library... ]  
  
#include <lber.h>  
#include <ldap.h>  
  
void ldap_control_free(LDAPControl *ctrl) ;  
void ldap_controls_free(LDAPControl *ctrls) ;
```

**DESCRIPTION** ldap\_controls\_free() and ldap\_control\_free() are routines which can be used to dispose of a single control or an array of controls allocated by other LDAP APIs.

**RETURN VALUES** None.

**ERRORS** No errors are defined for these functions.

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap\_error(3LDAP), ldap\_result(3LDAP), attributes(5)

**NAME** ldap\_delete, ldap\_delete\_s, ldap\_delete\_ext, ldap\_delete\_ext\_s – LDAP delete operation

**SYNOPSIS**

```
cc[ flag... ] file... -lldap[ library... ]

#include <lber.h>
#include <ldap.h>

int ldap_delete(LDAP *ld, char *dn);

int ldap_delete_s(LDAP *ld, char *dn);

int ldap_delete_ext(LDAP *ld, char *dn, LDAPControl **serverctrls,
LDAPControl **clientctrls, int *msgidp);

int ldap_delete_ext_s(LDAP *ld, char *dn, LDAPControl **serverctrls,
LDAPControl **clientctrls);
```

**DESCRIPTION**

The `ldap_delete_s()` function is used to perform an LDAP delete operation synchronously. It takes *dn*, the DN of the entry to be deleted. It returns an LDAP error code, indicating the success or failure of the operation.

The `ldap_delete()` function is used to perform an LDAP delete operation asynchronously. It takes the same parameters as `ldap_delete_s()`, but returns the message id of the request it initiated. The result of the delete can be obtained by a subsequent call to `ldap_result(3LDAP)`.

The `ldap_delete_ext()` function initiates an asynchronous delete operation and returns `LDAP_SUCCESS` if the request was successfully sent to the server, or else it returns a LDAP error code if not (see `ldap_error(3LDAP)`). If successful, `ldap_delete_ext()` places the message id of the request in *msgidp*. A subsequent call to `ldap_result()`, can be used to obtain the result of the add request.

The `ldap_delete_ext_s()` function initiates a synchronous delete operation and as such returns the result of the operation itself.

**ERRORS** `ldap_delete_s()` returns an LDAP error code which can be interpreted by calling one of `ldap_perror(3LDAP)` functions. `ldap_delete()` returns `-1` if something went wrong initiating the request. It returns the non-negative message id of the request if things were successful.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap(3LDAP)`, `ldap_error(3LDAP)`, `attributes(5)`

## ldap\_disptmpl(3LDAP)

<b>NAME</b>	ldap_disptmpl, ldap_init_templates, ldap_init_templates_buf, ldap_free_templates, ldap_first_disptmpl, ldap_next_disptmpl, ldap_oc2template, ldap_name2template, ldap_tmplattrs, ldap_first_tmplrow, ldap_next_tmplrow, ldap_first_tmplcol, ldap_next_tmplcol – LDAP display template functions
<b>SYNOPSIS</b>	<pre>cc[ <i>flag...</i> ] <i>file...</i> -lldap[ <i>library...</i> ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_init_templates(char *file, struct ldap_disptmpl **tmplist); int ldap_init_templates_buf(char *buf, unsigned long len, struct     ldap_disptmpl **tmplist); void ldap_free_templates(struct ldap_disptmpl *tmplist); struct ldap_disptmpl *ldap_first_disptmpl(struct ldap_disptmpl     *tmplist); struct ldap_disptmpl *ldap_next_disptmpl(struct ldap_disptmpl     *tmplist, struct ldap_disptmpl *tmpl); struct ldap_disptmpl *ldap_oc2template(char **oclist, struct     ldap_disptmpl *tmplist); struct ldap_disptmpl *ldap_name2template(char *name, struct     ldap_disptmpl *tmplist); char **ldap_tmplattrs(struct ldap_disptmpl *tmpl, char **includeatrs,     int exclude, unsigned long syntaxmask); struct ldap_tmplitem *ldap_first_tmplrow(struct ldap_disptmpl     *tmpl); struct ldap_tmplitem *ldap_next_tmplrow(struct ldap_disptmpl     *tmpl, struct ldap_tmplitem *row); struct ldap_tmplitem *ldap_first_tmplcol(struct ldap_disptmpl     *tmpl, struct ldap_tmplitem *row, struct ldap_tmplitem *col); struct ldap_tmplitem *ldap_next_tmplcol(struct ldap_disptmpl     *tmpl, struct ldap_tmplitem *row, struct ldap_tmplitem *col);</pre>
<b>DESCRIPTION</b>	<p>These functions provide a standard way to access LDAP entry display templates. Entry display templates provide a standard way for LDAP applications to display directory entries. The general idea is that it is possible to map the list of object class values present in an entry to an appropriate display template. Display templates are defined in a configuration file. See <code>ldaptemplates.conf(4)</code>. Each display template contains a pre-determined list of items, where each item generally corresponds to an attribute to be displayed. The items contain information and flags that the caller can use to display the attribute and values in a reasonable fashion. Each item has a syntaxid, which are described in the SYNTAX IDS section below. The <code>ldap_entry2text(3LDAP)</code> functions use the display template functions and produce text output.</p>

`ldap_init_templates()` reads a sequence of templates from a valid LDAP template configuration file (see `ldaptemplates.conf(4)`). Upon success, 0 is returned, and `tmplist` is set to point to a list of templates. Each member of the list is an `ldap_disptmpl` structure (defined below in the DISPTMPL Structure Elements section).

`ldap_init_templates_buf()` reads a sequence of templates from `buf` (whose size is `buflen`). `buf` should point to the data in the format defined for an LDAP template configuration file (see `ldaptemplates.conf(4)`). Upon success, 0 is returned, and `tmplist` is set to point to a list of templates.

The `LDAP_SET_DISPTMPL_APPDATA()` macro is used to set the value of the `dt_appdata` field in an `ldap_disptmpl` structure. This field is reserved for the calling application to use; it is not used internally.

The `LDAP_GET_DISPTMPL_APPDATA()` macro is used to retrieve the value in the `dt_appdata` field.

The `LDAP_IS_DISPTMPL_OPTION_SET()` macro is used to test a `ldap_disptmpl` structure for the existence of a template option. The options currently defined are: `LDAP_DTmpl_OPT_ADDABLE` (it is appropriate to allow entries of this type to be added), `LDAP_DTmpl_OPT_ALLOWMODRDN` (it is appropriate to offer the "modify rdn" operation), `LDAP_DTmpl_OPT_ALTVIEW` (this template is merely an alternate view of another template, typically used for templates pointed to be an `LDAP_SYN_LINKACTION` item).

`ldap_free_templates()` disposes of the templates allocated by `ldap_init_templates()`.

`ldap_first_disptmpl()` returns the first template in the list `tmplist`. The `tmplist` is typically obtained by calling `ldap_init_templates()`.

`ldap_next_disptmpl()` returns the template after `tmpl` in the template list `tmplist`. A NULL pointer is returned if `tmpl` is the last template in the list.

`ldap_oc2template()` searches `tmplist` for the best template to use to display an entry that has a specific set of `objectClass` values. `oclist` should be a null-terminated array of strings that contains the values of the `objectClass` attribute of the entry. A pointer to the first template where all of the object classes listed in one of the template's `dt_oclist` elements are contained in `oclist` is returned. A NULL pointer is returned if no appropriate template is found.

`ldap_tmplattrs()` returns a null-terminated array that contains the names of attributes that need to be retrieved if the template `tmpl` is to be used to display an entry. The attribute list should be freed using `ldap_value_free()`. The `includeattrs` parameter contains a null-terminated array of attributes that should always be included (it may be NULL if no extra attributes are required). If `syntaxmask` is non-zero, it is used to restrict the attribute set returned. If `exclude` is zero, only attributes where

## ldap\_disptmpl(3LDAP)

the logical AND of the template item syntax id and the *syntaxmask* is non-zero are included. If *exclude* is non-zero, attributes where the logical AND of the template item syntax id and the *syntaxmask* is non-zero are excluded.

`ldap_first_tmplrow()` returns a pointer to the first row of items in template *tmpl*.

`ldap_next_tmplrow()` returns a pointer to the row that follows *row* in template *tmpl*.

`ldap_first_tmplcol()` returns a pointer to the first item (in the first column) of row *row* within template *tmpl*. A pointer to an `ldap_tmplitem` structure (defined below in the TEMPLITEM Structure Elements section) is returned.

The `LDAP_SET_TEMPLITEM_APPDATA()` macro is used to set the value of the `ti_appdata` field in a `ldap_tmplitem` structure. This field is reserved for the calling application to use; it is not used internally.

The `LDAP_GET_TEMPLITEM_APPDATA()` macro is used to retrieve the value of the `ti_appdata` field.

The `LDAP_IS_TEMPLITEM_OPTION_SET()` macro is used to test a `ldap_tmplitem` structure for the existence of an item option. The options currently defined are: `LDAP_DITEM_OPT_READONLY` (this attribute should not be modified), `LDAP_DITEM_OPT_SORTVALUES` (it makes sense to sort the values), `LDAP_DITEM_OPT_SINGLEVALUED` (this attribute can only hold a single value), `LDAP_DITEM_OPT_VALUEREQUIRED` (this attribute must contain at least one value), `LDAP_DITEM_OPT_HIDEIFEMPTY` (do not show this item if there are no values), and `LDAP_DITEM_OPT_HIDEIFFALSE` (for boolean attributes only: hide this item if the value is FALSE).

`ldap_next_tmplcol()` returns a pointer to the item (column) that follows column *col* within row *row* of template *tmpl*.

### DISPTMPL Structure Elements

The `ldap_disptmpl` structure is defined as:

```
struct ldap_disptmpl {
    char                *dt_name;
    char                *dt_pluralname;
    char                *dt_iconname;
    unsigned long       dt_options;
    char                *dt_authattrname;
    char                *dt_defrdnattrname;
    char                *dt_defaddlocation;
    struct ldap_oclist  *dt_oclist;
    struct ldap_adddeflist *dt_adddeflist;
    struct ldap_tmplitem *dt_items;
    void                *dt_appdata;
    struct ldap_disptmpl *dt_next;
};
```



The `dt_name` member is the singular name of the template. The `dt_pluralname` is the plural name. The `dt_iconname` member will contain the name of an icon or other graphical element that can be used to depict entries that correspond to this display template. The `dt_options` contains options which may be tested using the `LDAP_IS_TMPLITEM_OPTION_SET()` macro.

The `dt_authattrname` contains the name of the DN-syntax attribute whose value(s) should be used to authenticate to make changes to an entry. If `dt_authattrname` is `NULL`, then authenticating as the entry itself is appropriate. The `dt_defrdnattrname` is the name of the attribute that is normally used to name entries of this type, for example, "cn" for person entries. The `dt_defaddlocation` is the distinguished name of an entry below which new entries of this type are typically created (its value is site-dependent).

`dt_oclist` is a pointer to a linked list of object class arrays, defined as:

```
struct ldap_oclist {
    char          **oc_objclasses;
    struct ldap_oclist *oc_next;
};
```

These are used by the `ldap_oc2template()` function.

`dt_adddeflist` is a pointer to a linked list of rules for defaulting the values of attributes when new entries are created. The `ldap_adddeflist` structure is defined as:

```
struct ldap_adddeflist {
    int          ad_source;
    char         *ad_attrname;
    char         *ad_value;
    struct ldap_adddeflist *ad_next;
};
```

The `ad_attrname` member contains the name of the attribute whose value this rule sets. If `ad_source` is `LDAP_ADSRC_CONSTANTVALUE` then the `ad_value` member contains the (constant) value to use. If `ad_source` is `LDAP_ADSRC_ADDERSDN` then `ad_value` is ignored and the distinguished name of the person who is adding the new entry is used as the default value for `ad_attrname`.

## TMPLITEM Structure Elements

The `ldap_tmplitem` structure is defined as:

```
struct ldap_tmplitem {
    unsigned long    ti_syntaxid;
    unsigned long    ti_options;
    char            *ti_attrname;
    char            *ti_label;
    char            **ti_args;
    struct ldap_tmplitem *ti_next_in_row;
    struct ldap_tmplitem *ti_next_in_col;
    void            *ti_appdata;
};
```

## ldap\_disptmpl(3LDAP)

**Syntax IDs** | Syntax ids are found in the `ldap_tmplitem` structure element `ti_syntaxid`, and they can be used to determine how to display the values for the attribute associated with an item. The `LDAP_GET_SYN_TYPE()` macro can be used to return a general type from a syntax id. The five general types currently defined are:  
`LDAP_SYN_TYPE_TEXT` (for attributes that are most appropriately shown as text),  
`LDAP_SYN_TYPE_IMAGE` (for JPEG or FAX format images),  
`LDAP_SYN_TYPE_BOOLEAN` (for boolean attributes), `LDAP_SYN_TYPE_BUTTON` (for attributes whose values are to be retrieved and display only upon request, for example, in response to the press of a button, a JPEG image is retrieved, decoded, and displayed), and `LDAP_SYN_TYPE_ACTION` (for special purpose actions such as "search for the entries where this entry is listed in the seeAlso attribute").

The `LDAP_GET_SYN_OPTIONS` macro can be used to retrieve an unsigned long bitmap that defines options. The only currently defined option is `LDAP_SYN_OPT_DEFER`, which (if set) implies that the values for the attribute should not be retrieved until requested.

There are sixteen distinct syntax ids currently defined. These generally correspond to one or more X.500 syntaxes.

`LDAP_SYN_CASEIGNORESTR` is used for text attributes which are simple strings whose case is ignored for comparison purposes.

`LDAP_SYN_MULTILINESTR` is used for text attributes which consist of multiple lines, for example, `postalAddress`, `homePostalAddress`, `multilineDescription`, or any attributes of syntax `caseIgnoreList`.

`LDAP_SYN_RFC822ADDR` is used for case ignore string attributes that are RFC-822 conformant mail addresses, for example, `mail`.

`LDAP_SYN_DN` is used for attributes with a Distinguished Name syntax, for example, `seeAlso`.

`LDAP_SYN_BOOLEAN` is used for attributes with a boolean syntax.

`LDAP_SYN_JPEGIMAGE` is used for attributes with a jpeg syntax, for example, `jpegPhoto`.

`LDAP_SYN_JPEGBUTTON` is used to provide a button (or equivalent interface element) that can be used to retrieve, decode, and display an attribute of jpeg syntax.

`LDAP_SYN_FAXIMAGE` is used for attributes with a photo syntax, for example, `Photo`. These are actually Group 3 Fax (T.4) format images.

`LDAP_SYN_FAXBUTTON` is used to provide a button (or equivalent interface element) that can be used to retrieve, decode, and display an attribute of photo syntax.

`LDAP_SYN_AUDIOBUTTON` is used to provide a button (or equivalent interface element) that can be used to retrieve and play an attribute of audio syntax. Audio values are in the "mu law" format, also known as "au" format.

LDAP\_SYN\_TIME is used for attributes with the UTCTime syntax, for example, lastModifiedTime. The value(s) should be displayed in complete date and time fashion.

LDAP\_SYN\_DATE is used for attributes with the UTCTime syntax, for example, lastModifiedTime. Only the date portion of the value(s) should be displayed.

LDAP\_SYN\_LABELLEDURL is used for labeledURL attributes.

LDAP\_SYN\_SEARCHACTION is used to define a search that is used to retrieve related information. If ti\_attrname is not NULL, it is assumed to be a boolean attribute which will cause no search to be performed if its value is FALSE. The ti\_args structure member will have four strings in it: ti\_args[ 0 ] should be the name of an attribute whose values are used to help construct a search filter or "-dn" is the distinguished name of the entry being displayed should be used, ti\_args[ 1 ] should be a filter pattern where any occurrences of "%v" are replaced with the value derived from ti\_args[ 0 ], ti\_args[ 2 ] should be the name of an additional attribute to retrieve when performing the search, and ti\_args[ 3 ] should be a human-consumable name for that attribute. The ti\_args[ 2 ] attribute is typically displayed along with a list of distinguished names when multiple entries are returned by the search.

LDAP\_SYN\_LINKACTION is used to define a link to another template by name. ti\_args[ 0 ] will contain the name of the display template to use. The ldap\_name2template() function can be used to obtain a pointer to the correct ldap\_disptmpl structure.

LDAP\_SYN\_ADDDNACTION and LDAP\_SYN\_VERIFYDNACTION are reserved as actions but currently undefined.

**ERRORS** The init template functions return LDAP\_TMPL\_ERR\_VERSION if *buf* points to data that is newer than can be handled, LDAP\_TMPL\_ERR\_MEM if there is a memory allocation problem, LDAP\_TMPL\_ERR\_SYNTAX if there is a problem with the format of the templates buffer or file. LDAP\_TMPL\_ERR\_FILE is returned by ldap\_init\_templates if the file cannot be read. Other functions generally return NULL upon error.

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_entry2text(3LDAP), ldaptemplates.conf(4), attributes(5)

## ldap\_entry2text(3LDAP)

<b>NAME</b>	ldap_entry2text, ldap_entry2text_search, ldap_entry2html, ldap_entry2html_search, ldap_vals2html, ldap_vals2text – LDAP entry display functions
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_entry2text(LDAP *ld, char *buf, LDAPMessage *entry, struct   ldap_disptmpl *tmpl, char **defattrs, char ***defvals, int   (*writeproc)(), void *writeparm, char *eol, int rdncount, unsigned   long opts);  int ldap_entry2text_search(LDAP *ld, char *dn, char *base,   LDAPMessage *entry, struct ldap_disptmpl *tmplist, char **defattrs,   char ***defvals, int (*writeproc)(), void *writeparm, char *eol, int   rdncount, unsigned long opts);  int ldap_vals2text(LDAP *ld, char *buf, char **vals, char *label, int   labelwidth, unsigned long syntaxid, int (*writeproc)(), void *writeparm,   char *eol, int rdncount);  int ldap_entry2html(LDAP *ld, char *buf, LDAPMessage *entry, struct   ldap_disptmpl *tmpl, char **defattrs, char ***defvals, int   (*writeproc)(), void *writeparm, char *eol, int rdncount, unsigned   long opts, char *urlprefix, char *base);  int ldap_entry2html_search(LDAP *ld, char *dn, LDAPMessage *entry,   struct ldap_disptmpl *tmplist, char **defattrs, char ***defvals, int   (*writeproc)(), void *writeparm, char *eol, int rdncount, unsigned   long opts, char *urlprefix);  int ldap_vals2html(LDAP *ld, char *buf, char **vals, char *label, int   labelwidth, unsigned long syntaxid, int (*writeproc)(), void   *writeparm, char *eol, int rdncount, char *urlprefix);  #define LDAP_DISP_OPT_AUTOLABELWIDTH 0x00000001  #define LDAP_DISP_OPT_HTMLBODYONLY 0x00000002  #define LDAP_DTmpl_BUFSIZ 2048</pre>
<b>DESCRIPTION</b>	<p>These functions use the LDAP display template functions (see <code>ldap_disptmpl(3LDAP)</code> and <code>ldap_templates.conf(4)</code>) to produce a plain text or an HyperText Markup Language (HTML) display of an entry or a set of values. Typical plain text output produced for an entry might look like:</p> <pre>"Barbara J Jensen, Information Technology Division" Also Known As: Babs Jensen Barbara Jensen Barbara J Jensen E-Mail Address: bjensen@terminator.rs.itd.umich.edu Work Address:</pre>

```

535 W. William
Ann Arbor, MI 48103
Title:
Mythical Manager, Research Systems
...

```

The exact output produced will depend on the display template configuration. HTML output is similar to the plain text output, but more richly formatted.

`ldap_entry2text()` produces a text representation of *entry* and writes the text by calling the *writeproc* function. All of the attributes values to be displayed must be present in *entry*; no interaction with the LDAP server will be performed within `ldap_entry2text`. `ld` is the LDAP pointer obtained by a previous call to `ldap_open`. *writeproc* should be declared as:

```

int writeproc( writeparm, p, len )
void *writeparm;
char *p;
int len;

```

where *p* is a pointer to text to be written and *len* is the length of the text. *p* is guaranteed to be zero-terminated. Lines of text are terminated with the string *col*. *buf* is a pointer to a buffer of size `LDAP_DTmpl_BUFSIZ` or larger. If *buf* is NULL then a buffer is allocated and freed internally. *tmpl* is a pointer to the display template to be used (usually obtained by calling `ldap_oc2template`). If *tmpl* is NULL, no template is used and a generic display is produced. *defattrs* is a NULL-terminated array of LDAP attribute names which you wish to provide default values for (only used if *entry* contains no values for the attribute). An array of NULL-terminated arrays of default values corresponding to the attributes should be passed in *defvals*. The *rdncount* parameter is used to limit the number of Distinguished Name (DN) components that are actually displayed for DN attributes. If *rdncount* is zero, all components are shown. *opts* is used to specify output options. The only values currently allowed are zero (default output), `LDAP_DISP_OPT_AUTOLABELWIDTH` which causes the width for labels to be determined based on the longest label in *tmpl*, and `LDAP_DISP_OPT_HTMLBODYONLY`. The `LDAP_DISP_OPT_HTMLBODYONLY` option instructs the library not to include `<HTML>`, `<HEAD>`, `<TITLE>`, and `<BODY>` tags. In other words, an HTML fragment is generated, and the caller is responsible for prepending and appending the appropriate HTML tags to construct a correct HTML document.

`ldap_entry2text_search()` is similar to `ldap_entry2text`, and all of the like-named parameters have the same meaning except as noted below. If *base* is not NULL, it is the search base to use when executing search actions. If it is NULL, search action template items are ignored. If *entry* is not NULL, it should contain the *objectClass* attribute values for the entry to be displayed. If *entry* is NULL, *dn* must not be NULL, and `ldap_entry2text_search` will retrieve the *objectClass* values itself by calling `ldap_search_s`. `ldap_entry2text_search` will determine the appropriate display template to use by calling `ldap_oc2template`, and will call `ldap_search_s` to retrieve any attribute values to be displayed. The *tmplist* parameter is a pointer to the entire list of templates available (usually obtained by

## ldap\_entry2text(3LDAP)

calling `ldap_init_templates` or `ldap_init_templates_buf`). If `tmplist` is NULL, `ldap_entry2text_search` will attempt to read a load templates from the default template configuration file `ETCDIR/ldaptemplates.conf`.

`ldap_vals2text` produces a text representation of a single set of LDAP attribute values. The `ld`, `buf`, `writeproc`, `writeparm`, `eol`, and `rdncount` parameters are the same as the like-named parameters for `ldap_entry2text`. `vals` is a NULL-terminated list of values, usually obtained by a call to `ldap_get_values`. `label` is a string shown next to the values (usually a friendly form of an LDAP attribute name). `labelwidth` specifies the label margin, which is the number of blank spaces displayed to the left of the values. If zero is passed, a default label width is used. `syntaxid` is a display template attribute syntax identifier (see `ldap_disptmpl(3LDAP)` for a list of the pre-defined LDAP\_SYN... values).

`ldap_entry2html` produces an HTML representation of `entry`. It behaves exactly like `ldap_entry2text(3LDAP)`, except for the formatted output and the addition of two parameters. `urlprefix` is the starting text to use when constructing an LDAP URL. The default is the string `ldap:///`. The second additional parameter, `base`, the search base to use when executing search actions. If it is NULL, search action template items are ignored.

`ldap_entry2html_search` behaves exactly like `ldap_entry2text_search(3LDAP)`, except HTML output is produced and one additional parameter is required. `urlprefix` is the starting text to use when constructing an LDAP URL. The default is the string `ldap:///`

`ldap_vals2html` behaves exactly like `ldap_vals2text`, except HTML output is produced and one additional parameter is required. `urlprefix` is the starting text to use when constructing an LDAP URL. The default is the string `ldap:///`

**ERRORS** These functions all return an LDAP error code. LDAP\_SUCCESS is returned if no error occurs. See `ldap_error(3LDAP)` for details. The `ld_errno` field of the `ld` parameter is also set to indicate the error.

**FILES** `ETCDIR/ldaptemplates.conf`

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap(3LDAP)`, `ldap_disptmpl(3LDAP)`, `ldaptemplates.conf(4)`, `attributes(5)`

<b>NAME</b>	ldap_error, ldap_err2string, ldap_perror, ldap_result2error – LDAP protocol error handling functions																
<b>SYNOPSIS</b>	<pre>cc[ flag... ] file... -lldap[ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  char *ldap_err2string(int err);  void ldap_perror(LDAP *ld, const char *s);  int ldap_result2error(LDAP *ld, LDAPMessage *res, int freeit);</pre>																
<b>DESCRIPTION</b>	<p>These functions interpret the error codes that are returned by the LDAP API routines. The <code>ldap_perror()</code> and <code>ldap_result2error()</code> functions are deprecated for all new development. Use <code>ldap_err2string()</code> instead.</p> <p>You can also use <code>ldap_parse_sasl_bind_result(3LDAP)</code>, <code>ldap_parse_extended_result(3LDAP)</code>, and <code>ldap_parse_result(3LDAP)</code> to provide error handling and interpret error codes returned by LDAP API functions.</p> <p>The <code>ldap_err2string()</code> function takes <i>err</i>, a numeric LDAP error code, returned either by <code>ldap_parse_result(3LDAP)</code> or another LDAP API call. It returns an informative, null-terminated, character string that describes the error.</p> <p>The <code>ldap_result2error()</code> function takes <i>res</i>, a result produced by <code>ldap_result(3LDAP)</code> or other synchronous LDAP calls, and returns the corresponding error code. If the <i>freeit</i> parameter is non-zero, it indicates that the <i>res</i> parameter should be freed by a call to <code>ldap_msgfree(3LDAP)</code> after the error code has been extracted.</p> <p>Similar to the way <code>perror(3C)</code> works, the <code>ldap_perror()</code> function can be called to print an indication of the error to standard error.</p>																
<b>ERRORS</b>	<p>The possible values for an LDAP error code are:</p> <table border="0"> <tr> <td style="padding-right: 20px;">LDAP_SUCCESS</td> <td>The request was successful.</td> </tr> <tr> <td>LDAP_OPERATIONS_ERROR</td> <td>An operations error occurred.</td> </tr> <tr> <td>LDAP_PROTOCOL_ERROR</td> <td>A protocol violation was detected.</td> </tr> <tr> <td>LDAP_TIMELIMIT_EXCEEDED</td> <td>An LDAP time limit was exceeded.</td> </tr> <tr> <td>LDAP_SIZELIMIT_EXCEEDED</td> <td>An LDAP size limit was exceeded.</td> </tr> <tr> <td>LDAP_COMPARE_FALSE</td> <td>A compare operation returned false.</td> </tr> <tr> <td>LDAP_COMPARE_TRUE</td> <td>A compare operation returned true.</td> </tr> <tr> <td>LDAP_STRONG_AUTH_NOT_SUPPORTED</td> <td>The LDAP server does not support strong authentication.</td> </tr> </table>	LDAP_SUCCESS	The request was successful.	LDAP_OPERATIONS_ERROR	An operations error occurred.	LDAP_PROTOCOL_ERROR	A protocol violation was detected.	LDAP_TIMELIMIT_EXCEEDED	An LDAP time limit was exceeded.	LDAP_SIZELIMIT_EXCEEDED	An LDAP size limit was exceeded.	LDAP_COMPARE_FALSE	A compare operation returned false.	LDAP_COMPARE_TRUE	A compare operation returned true.	LDAP_STRONG_AUTH_NOT_SUPPORTED	The LDAP server does not support strong authentication.
LDAP_SUCCESS	The request was successful.																
LDAP_OPERATIONS_ERROR	An operations error occurred.																
LDAP_PROTOCOL_ERROR	A protocol violation was detected.																
LDAP_TIMELIMIT_EXCEEDED	An LDAP time limit was exceeded.																
LDAP_SIZELIMIT_EXCEEDED	An LDAP size limit was exceeded.																
LDAP_COMPARE_FALSE	A compare operation returned false.																
LDAP_COMPARE_TRUE	A compare operation returned true.																
LDAP_STRONG_AUTH_NOT_SUPPORTED	The LDAP server does not support strong authentication.																

## ldap\_error(3LDAP)

LDAP_STRONG_AUTH_REQUIRED	Strong authentication is required for the operation.
LDAP_PARTIAL_RESULTS	Only partial results are returned.
LDAP_NO_SUCH_ATTRIBUTE	The attribute type specified does not exist in the entry.
LDAP_UNDEFINED_TYPE	The attribute type specified is invalid.
LDAP_INAPPROPRIATE_MATCHING	The filter type is not supported for the specified attribute.
LDAP_CONSTRAINT_VIOLATION	An attribute value specified violates some constraint. For example, a <code>postalAddress</code> has too many lines, or a line that is too long.
LDAP_TYPE_OR_VALUE_EXISTS	An attribute type or attribute value specified already exists in the entry.
LDAP_INVALID_SYNTAX	An invalid attribute value was specified.
LDAP_NO_SUCH_OBJECT	The specified object does not exist in the directory.
LDAP_ALIAS_PROBLEM	An alias in the directory points to a nonexistent entry.
LDAP_INVALID_DN_SYNTAX	A syntactically invalid DN was specified.
LDAP_IS_LEAF	The object specified is a leaf.
LDAP_ALIAS_DEREF_PROBLEM	A problem was encountered when dereferencing an alias.
LDAP_INAPPROPRIATE_AUTH	Inappropriate authentication was specified. For example, <code>LDAP_AUTH_SIMPLE</code> was specified and the entry does not have a <code>userPassword</code> attribute.
LDAP_INVALID_CREDENTIALS	Invalid credentials were presented, for example, the wrong password.
LDAP_INSUFFICIENT_ACCESS	The user has insufficient access to perform the operation.
LDAP_BUSY	The DSA is busy.
LDAP_UNAVAILABLE	The DSA is unavailable.
LDAP_UNWILLING_TO_PERFORM	The DSA is unwilling to perform the operation.
LDAP_LOOP_DETECT	A loop was detected.
LDAP_NAMING_VIOLATION	A naming violation occurred.



## ldap\_error(3LDAP)

LDAP_OBJECT_CLASS_VIOLATION	An object class violation occurred. For example, a must attribute was missing from the entry.
LDAP_NOT_ALLOWED_ON_NONLEAF	The operation is not allowed on a nonleaf object.
LDAP_NOT_ALLOWED_ON_RDN	The operation is not allowed on an RDN.
LDAP_ALREADY_EXISTS	The entry already exists.
LDAP_NO_OBJECT_CLASS_MODS	Object class modifications are not allowed.
LDAP_OTHER	An unknown error occurred.
LDAP_SERVER_DOWN	The LDAP library cannot contact the LDAP server.
LDAP_LOCAL_ERROR	Some local error occurred. This is usually a failed <code>malloc()</code> .
LDAP_ENCODING_ERROR	An error was encountered encoding parameters to send to the LDAP server.
LDAP_DECODING_ERROR	An error was encountered decoding a result from the LDAP server.
LDAP_TIMEOUT	A time limit was exceeded while waiting for a result.
LDAP_AUTH_UNKNOWN	The authentication method specified to <code>ldap_bind(3LDAP)</code> is not known.
LDAP_FILTER_ERROR	An invalid filter was supplied to <code>ldap_search(3LDAP)</code> , for example, unbalanced parentheses.
LDAP_PARAM_ERROR	An LDAP function was called with a bad parameter, for example, a NULL <i>ld</i> pointer, and the like.
LDAP_NO_MEMORY	A memory allocation call failed in an LDAP library function, for example, <code>malloc(3C)</code> .
LDAP_CONNECT_ERROR	The LDAP client has either lost its connection to an LDAP server or it cannot establish a connection.
LDAP_NOT_SUPPORTED	The requested functionality is not supported., for example, when an LDAPv2 client requests some LDAPv3 functionality.
LDAP_CONTROL_NOT_FOUND	An LDAP client requested a control not found in the list of supported controls sent by the server.

## ldap\_error(3LDAP)

LDAP_NO_RESULTS_RETURNED	The LDAP server sent no results.
LDAP_MORE_RESULTS_TO_RETURN	More results are chained in the message chain.
LDAP_CLIENT_LOOP	A loop has been detected, for example, when following referrals.
LDAP_REFERRAL_LIMIT_EXCEEDED	The referral exceeds the hop limit. The hop limit determines the number of servers that the client can hop through to retrieve data.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap(3LDAP)`, `ldap_bind(3LDAP)`, `ldap_msgfree(3LDAP)`, `ldap_parse_extended_result(3LDAP)`, `ldap_parse_result(3LDAP)`, `ldap_parse_sasl_bind_result(3LDAP)`, `ldap_search(3LDAP)`, `malloc(3C)`, `perror(3C)`, `attributes(5)`

**NAME** ldap\_first\_attribute, ldap\_next\_attribute – step through LDAP entry attributes

**SYNOPSIS**

```
cc [ flag... ] file... -lldap[ library... ]

#include <lber.h>
#include <ldap.h>

char *ldap_first_attribute(LDAP *ld, LDAPMessage *entry, BerElement
**berptr) ;

char *ldap_next_attribute(LDAP *ld, LDAPMessage *entry, BerElement
*ber) ;
```

**DESCRIPTION**

The `ldap_first_attribute()` and `ldap_next_attribute()` functions are used to step through the attributes in an LDAP entry. `ldap_first_attribute()` takes an *entry* as returned by `ldap_first_entry(3LDAP)` or `ldap_next_entry(3LDAP)` and returns a pointer to memory that must be freed with a call to `ldap_memfree(3LDAP)`.

They also return, in *berptr*, a pointer to a `BerElement` it has allocated to keep track of its current position. This pointer should be passed to subsequent calls to `ldap_next_attribute()` and is used to effectively step through the entry's attributes. This pointer is freed by `ldap_next_attribute()` when there are no more attributes, that is, when `ldap_next_attribute()` returns `NULL`. Otherwise, the caller is responsible for freeing the `BerElement` pointed to by *berptr* when it is no longer needed by calling `ber_free(3LDAP)`. When calling `ber_free(3LDAP)` in this instance, be sure the second argument is '0'.

The attribute names returned are suitable for inclusion in a call to `ldap_get_values(3LDAP)` to retrieve the attribute's values.

**ERRORS**

If an error occurs, `NULL` is returned and the `ld_errno` field in the *ld* parameter is set to indicate the error. See `ldap_error(3LDAP)` for a description of possible error codes.

**ATTRIBUTES**

See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap(3LDAP)`, `ldap_first_entry(3LDAP)`, `ldap_get_values(3LDAP)`, `ldap_error(3LDAP)`, `attributes(5)`

**NOTES**

The `ldap_first_attribute()` function allocates memory that might need to be freed by the caller by means of `ber_free(3LDAP)`.

## ldap\_first\_entry(3LDAP)

<b>NAME</b>	ldap_first_entry, ldap_next_entry, ldap_count_entries, ldap_count_references, ldap_first_reference, ldap_next_reference – LDAP entry parsing and counting functions
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  LDAPMessage *ldap_first_entry(LDAP*ld, LDAPMessage *result); LDAPMessage *ldap_next_entry(LDAP *ld, LDAPMessage *entry); ldap_count_entries(LDAP *ld, LDAPMessage *result); LDAPMessage *ldap_first_reference(LDAP *ld, LDAPMessage *res); LDAPMessage *ldap_next_reference(LDAP *ld, LDAPMessage *res); int ldap_count_references(LDAP *ld, LDAPMessage *res);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to parse results received from <code>ldap_result(3LDAP)</code> or the synchronous LDAP search operation functions <code>ldap_search_s(3LDAP)</code> and <code>ldap_search_st(3LDAP)</code>.</p> <p>The <code>ldap_first_entry()</code> function is used to retrieve the first entry in a chain of search results. It takes the <i>result</i> as returned by a call to <code>ldap_result(3LDAP)</code> or <code>ldap_search_s(3LDAP)</code> or <code>ldap_search_st(3LDAP)</code> and returns a pointer to the first entry in the result.</p> <p>This pointer should be supplied on a subsequent call to <code>ldap_next_entry()</code> to get the next entry, the result of which should be supplied to the next call to <code>ldap_next_entry()</code>, etc. <code>ldap_next_entry()</code> will return <code>NULL</code> when there are no more entries. The entries returned from these calls are used in calls to the functions described in <code>ldap_get_dn(3LDAP)</code>, <code>ldap_first_attribute(3LDAP)</code>, <code>ldap_get_values(3LDAP)</code>, etc.</p> <p>A count of the number of entries in the search result can be obtained by calling <code>ldap_count_entries()</code>.</p> <p><code>ldap_first_reference()</code> and <code>ldap_next_reference()</code> are used to step through and retrieve the list of continuation references from a search result chain.</p> <p>The <code>ldap_count_references()</code> function is used to count the number of references that are contained in and remain in a search result chain.</p>
<b>ERRORS</b>	If an error occurs in <code>ldap_first_entry()</code> or <code>ldap_next_entry()</code> , <code>NULL</code> is returned and the <code>ld_errno</code> field in the <i>ld</i> parameter is set to indicate the error. If an error occurs in <code>ldap_count_entries()</code> , <code>-1</code> is returned, and <code>ld_errno</code> is set appropriately. See <code>ldap_error(3LDAP)</code> for a description of possible error codes.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for a description of the following attributes:

ldap\_first\_entry(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_result(3LDAP), ldap\_search(3LDAP),  
ldap\_first\_attribute(3LDAP), ldap\_get\_values(3LDAP),  
ldap\_get\_dn(3LDAP), attributes(5)

## ldap\_first\_message(3LDAP)

<b>NAME</b>	ldap_first_message, ldap_count_messages, ldap_next_message, ldap_msgtype – LDAP message processing functions						
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_count_messages(LDAP *ld, LDAPMessage *res);  LDAPMessage *ldap_first_message(LDAP *ld, LDAPMessage *res);  LDAPMessage *ldap_next_message(LDAP *ld, LDAPMessage *msg);  int ldap_msgtype(LDAPMessage *res);</pre>						
<b>DESCRIPTION</b>	<p>ldap_count_messages() is used to count the number of messages that remain in a chain of results if called with a message, entry, or reference returned by ldap_first_message(), ldap_next_message(), ldap_first_entry(), ldap_next_entry(), ldap_first_reference(), and ldap_next_reference().</p> <p>ldap_first_message() and ldap_next_message() functions are used to step through the list of messages in a result chain returned by ldap_result().</p> <p>ldap_msgtype() function returns the type of an LDAP message.</p>						
<b>RETURN VALUES</b>	<p>ldap_first_message() and ldap_next_message() return LDAPMessage which can include referral messages, entry messages and result messages.</p> <p>ldap_count_messages() returns the number of messages contained in a chain of results.</p>						
<b>ERRORS</b>	ldap_first_message() and ldap_next_message() return NULL when no more messages exist. NULL is also returned if an error occurs while stepping through the entries, in which case the error parameters in the session handle ld will be set to indicate the error.						
<b>ATTRIBUTES</b>	See attributes(5) for a description of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWcsl (32-bit) SUNWcslx (64-bit)</td></tr><tr><td>Interface Stability</td><td>Evolving</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)						
Interface Stability	Evolving						
<b>SEE ALSO</b>	ldap_error(3LDAP), ldap_result(3LDAP), attributes(5)						

<b>NAME</b>	ldap_friendly, ldap_friendly_name, ldap_free_friendlymap – LDAP attribute remapping functions						
<b>SYNOPSIS</b>	<pre>cc[ flag... ] file... -lldap[ library... ] #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  char *ldap_friendly_name(char *filename, char *name, FriendlyMap **map) ;  void ldap_free_friendlymap(FriendlyMap **map) ;</pre>						
<b>DESCRIPTION</b>	<p>This function is used to map one set of strings to another. Typically, this is done for country names, to map from the two-letter country codes to longer more readable names. The mechanism is general enough to be used with other things, though.</p> <p><i>filename</i> is the name of a file containing the unfriendly to friendly mapping, <i>name</i> is the unfriendly name to map to a friendly name, and <i>map</i> is a result-parameter that should be set to NULL on the first call. It is then used to hold the mapping in core so that the file need not be read on subsequent calls.</p> <p>For example:</p> <pre> FriendlyMap *map = NULL; printf( "unfriendly %s =&gt; friendly %s\n", name, ldap_friendly_name( "ETCDIR/ldapfriendly", name, &amp;map ) );</pre> <p>The mapping file should contain lines like this: unfriendlyname\tfriendlyname. Lines that begin with a '#' character are comments and are ignored.</p> <p>The ldap_free_friendlymap() call is used to free structures allocated by ldap_friendly_name() when no more calls to ldap_friendly_name() are to be made.</p>						
<b>ERRORS</b>	NULL is returned by ldap_friendly_name() if there is an error opening <i>filename</i> , or if the file has a bad format, or if the <i>map</i> parameter is NULL.						
<b>FILES</b>	ETCDIR/ldapfriendly.conf						
<b>ATTRIBUTES</b>	See attributes(5) for a description of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWcsl (32-bit) SUNWcslx (64-bit)</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)						
Interface Stability	Evolving						
<b>SEE ALSO</b>	ldap(3LDAP), attributes(5)						

## ldap\_get\_dn(3LDAP)

<b>NAME</b>	ldap_get_dn, ldap_explode_dn, ldap_dn2ufn, ldap_is_dns_dn, ldap_explode_dns, ldap_dns_to_dn – LDAP DN handling functions
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  char *ldap_get_dn(LDAP *ld, LDAPMessage *entry); char **ldap_explode_dn(char *dn, int notypes); char *ldap_dn2ufn(char *dn); int ldap_is_dns_dn(char *dn); char **ldap_explode_dns(char *dn); char *ldap_dns_to_dn(char *dns_name, int *nameparts);</pre>
<b>DESCRIPTION</b>	<p>These functions allow LDAP entry names (Distinguished Names, or DN's) to be obtained, parsed, converted to a user-friendly form, and tested. A DN has the form described in RFC 1779 <i>A String Representation of Distinguished Names</i>, unless it is an experimental DNS-style DN which takes the form of an RFC 822 mail address.</p> <p>The <code>ldap_get_dn()</code> function takes an <i>entry</i> as returned by <code>ldap_first_entry(3LDAP)</code> or <code>ldap_next_entry(3LDAP)</code> and returns a copy of the entry's DN. Space for the DN will have been obtained by means of <code>malloc(3C)</code>, and should be freed by the caller by a call to <code>free(3C)</code>.</p> <p>The <code>ldap_explode_dn()</code> function takes a DN as returned by <code>ldap_get_dn()</code> and breaks it up into its component parts. Each part is known as a Relative Distinguished Name, or RDN. <code>ldap_explode_dn()</code> returns a null-terminated array, each component of which contains an RDN from the DN. The <i>notypes</i> parameter is used to request that only the RDN values be returned, not their types. For example, the DN "cn=Bob,c=US" would return as either { "cn=Bob", "c=US", NULL } or { "Bob", "US", NULL }, depending on whether <i>notypes</i> was 0 or 1, respectively. The result can be freed by calling <code>ldap_value_free(3LDAP)</code>.</p> <p><code>ldap_dn2ufn()</code> is used to turn a DN as returned by <code>ldap_get_dn()</code> into a more user-friendly form, stripping off type names. See RFC 1781 "Using the Directory to Achieve User Friendly Naming" for more details on the UFN format. The space for the UFN returned is obtained by a call to <code>malloc(3C)</code>, and the user is responsible for freeing it by means of a call to <code>free(3C)</code>.</p> <p><code>ldap_is_dns_dn()</code> returns non-zero if the dn string is an experimental DNS-style DN (generally in the form of an RFC 822 e-mail address). It returns zero if the dn appears to be an RFC 1779 format DN.</p> <p><code>ldap_explode_dns()</code> takes a DNS-style DN and breaks it up into its component parts. <code>ldap_explode_dns()</code> returns a null-terminated array. For example, the DN "mcs.umich.edu" will return { "mcs", "umich", "edu", NULL }. The result can be freed by calling <code>ldap_value_free(3LDAP)</code>.</p>



ldap\_get\_dn(3LDAP)

ldap\_dns\_to\_dn() converts a DNS domain name into an X.500 distinguished name. A string distinguished name and the number of nameparts is returned.

**ERRORS** If an error occurs in ldap\_get\_dn(), NULL is returned and the ld\_errno field in the ld parameter is set to indicate the error. See ldap\_error(3LDAP) for a description of possible error codes. ldap\_explode\_dn(), ldap\_explode\_dns() and ldap\_dn2ufn() will return NULL with errno(3C) set appropriately in case of trouble.

If an error in ldap\_dns\_to\_dn() is encountered zero is returned. The caller should free the returned string if it is non-zero.

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_first\_entry(3LDAP), ldap\_error(3LDAP), ldap\_value\_free(3LDAP)

**NOTES** These functions allocate memory that the caller must free.

## ldap\_get\_entry\_controls(3LDAP)

<b>NAME</b>	ldap_get_entry_controls – get the LDAP controls included with a directory entry in a set of search results								
<b>SYNOPSIS</b>	<pre>cc -flag ... file ...-lldap [-library ...] #include &lt;ldap.h&gt;  int ldap_get_entry_controls(LDAP *ld, LDAPMessage *entry,     LDAPControl ***serverctrlsp);</pre>								
<b>DESCRIPTION</b>	<p>The <code>ldap_get_entry_controls()</code> function retrieves the LDAP v3 controls included in a directory entry in a chain of search results. The LDAP controls are specified in an array of <code>LDAPControl</code> structures. Each <code>LDAPControl</code> structure represents an LDAP control. The function takes <i>entry</i> as a parameter, which points to an <code>LDAPMessage</code> structure that represents an entry in a chain of search results.</p> <p>The entry notification controls that are used with persistent search controls are the only controls that are returned with individual entries. Other controls are returned with results sent from the server. You can call <code>ldap_parse_result()</code> to retrieve those controls.</p>								
<b>ERRORS</b>	<p><code>ldap_get_entry_controls()</code> returns the following error codes.</p> <table><tr><td><code>LDAP_SUCCESS</code></td><td>LDAP controls were successfully retrieved.</td></tr><tr><td><code>LDAP_DECODING_ERROR</code></td><td>An error occurred when decoding the BER-encoded message.</td></tr><tr><td><code>LDAP_PARAM_ERROR</code></td><td>An invalid parameter was passed to the function.</td></tr><tr><td><code>LDAP_NO_MEMORY</code></td><td>Memory cannot be allocated.</td></tr></table>	<code>LDAP_SUCCESS</code>	LDAP controls were successfully retrieved.	<code>LDAP_DECODING_ERROR</code>	An error occurred when decoding the BER-encoded message.	<code>LDAP_PARAM_ERROR</code>	An invalid parameter was passed to the function.	<code>LDAP_NO_MEMORY</code>	Memory cannot be allocated.
<code>LDAP_SUCCESS</code>	LDAP controls were successfully retrieved.								
<code>LDAP_DECODING_ERROR</code>	An error occurred when decoding the BER-encoded message.								
<code>LDAP_PARAM_ERROR</code>	An invalid parameter was passed to the function.								
<code>LDAP_NO_MEMORY</code>	Memory cannot be allocated.								
<b>ATTRIBUTES</b>	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWcsl (32-bit)</td></tr><tr><td></td><td>SUNWcslx (64-bit)</td></tr><tr><td>Interface Stability</td><td>Evolving</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsl (32-bit)		SUNWcslx (64-bit)	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWcsl (32-bit)								
	SUNWcslx (64-bit)								
Interface Stability	Evolving								
<b>SEE ALSO</b>	<code>ldap_error(3LDAP)</code> , <code>ldap_parse_result(3LDAP)</code> , <code>attributes(5)</code>								

<b>NAME</b>	ldap_getfilter, ldap_init_getfilter, ldap_init_getfilter_buf, ldap_getfilter_free, ldap_getfirstfilter, ldap_getnextfilter, ldap_setfilteraffixes, ldap_build_filter – LDAP filter generating functions
<b>SYNOPSIS</b>	<pre>cc[ <i>flag...</i> ] <i>file...</i> -lldap[ <i>library...</i> ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt; #define LDAP_FILT_MAXSIZ    1024  LDAPFiltDesc *ldap_init_getfilter(char *file);  LDAPFiltDesc *ldap_init_getfilter_buf(char *buf, long buflen);  ldap_getfilter_free(LDAFFiltDesc *lfdp);  LDAPFiltInfo *ldap_getfirstfilter(LDAFFiltDesc *lfdp, char *tagpat,     char *value);  LDAPFiltInfo *ldap_getnextfilter(LDAFFiltDesc *lfdp);  void ldap_setfilteraffixes(LDAFFiltDesc *lfdp, char *prefix, char     *suffix);  void ldap_build_filter(char *buf, unsigned long buflen, char *pattern,     char *prefix, char *suffix, char *attr, char *value, char **valwords);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to generate filters to be used in ldap_search(3LDAP) or ldap_search_s(3LDAP). Either ldap_init_getfilter or ldap_init_getfilter_buf must be called prior to calling any of the other functions except ldap_build_filter.</p> <p>ldap_init_getfilter() takes a file name as its only argument. The contents of the file must be a valid LDAP filter configuration file (see ldapfilter.conf(4)). If the file is successfully read, a pointer to an LDAFFiltDesc is returned. This is an opaque object that is passed in subsequent get filter calls.</p> <p>ldap_init_getfilter_buf() reads from <i>buf</i>, whose length is <i>buflen</i>, the LDAP filter configuration information. <i>buf</i> must point to the contents of a valid LDAP filter configuration file. See ldapfilter.conf(4). If the filter configuration information is successfully read, a pointer to an LDAFFiltDesc is returned. This is an opaque object that is passed in subsequent get filter calls.</p> <p>ldap_getfilter_free() deallocates the memory consumed by ldap_init_getfilter. Once it is called, the LDAFFiltDesc is no longer valid and cannot be used again.</p> <p>ldap_getfirstfilter() retrieves the first filter that is appropriate for <i>value</i>. Only filter sets that have tags that match the regular expression <i>tagpat</i> are considered. ldap_getfirstfilter returns a pointer to an LDAFFiltInfo structure, which contains a filter with <i>value</i> inserted as appropriate in <i>lfi_filter</i>, a text match description in <i>lfi_desc</i>, <i>lfi_scope</i> set to indicate the search scope, and</p>

## ldap\_getfilter(3LDAP)

`lfi_isexact` set to indicate the type of filter. `NULL` is returned if no matching filters are found. `lfi_scope` will be one of `LDAP_SCOPE_BASE`, `LDAP_SCOPE_ONELEVEL`, or `LDAP_SCOPE_SUBTREE`. `lfi_isexact` will be zero if the filter has any '~' or '\*' characters in it and non-zero otherwise.

`ldap_getnextfilter()` retrieves the next appropriate filter in the filter set that was determined when `ldap_getfirstfilter` was called. It returns `NULL` when the list has been exhausted.

`ldap_setfilteraffixes()` sets a *prefix* to be prepended and a *suffix* to be appended to all filters returned in the future.

`ldap_build_filter()` constructs an LDAP search filter in *buf*. *buflen* is the size, in bytes, of the largest filter *buf* can hold. A pattern for the desired filter is passed in *pattern*. Where the string %a appears in the pattern it is replaced with *attr*. *prefix* is pre-pended to the resulting filter, and *suffix* is appended. Either can be `NULL`, in which case they are not used. *value* and *valwords* are used when the string %v appears in *pattern*. See `ldapfilter.conf(4)` for a description of how %v is handled.

**ERRORS** `NULL` is returned by `ldap_init_getfilter` if there is an error reading *file*. `NULL` is returned by `ldap_getfirstfilter` and `ldap_getnextfilter` when there are no more appropriate filters to return.

**FILES** `ETCDIR/ldapfilter.conf` LDAP filtering routine configuration file.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap(3LDAP)`, `ldapfilter.conf(4)`, `attributes(5)`

**NOTES** The return values for all of these functions are declared in the `<ldap.h>` header file. Some functions may allocate memory which must be freed by the calling application.

<b>NAME</b>	ldap_get_lang_values, ldap_get_lang_values_len – return an attribute’s values that matches a specified language subtype
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lldap [-library ...] #include &lt;ldap.h&gt;  char **ldap_get_lang_values(LDAP *ld, LDAPMessage *entry, const     char *target, char **type);  struct berval **ldap_get_lang_values_len(LDAP *ld, LDAPMessage     *entry, const char *target, char **type);</pre>
<b>DESCRIPTION</b>	<p>The <code>ldap_get_lang_values()</code> function returns an array of an attribute’s string values that matches a specified language subtype. To retrieve the binary data from an attribute, call the <code>ldap_get_lang_values_len()</code> function instead.</p> <p><code>ldap_get_lang_values()</code> should be called to retrieve a null-terminated array of an attribute’s string values that match a specified language subtype. The <i>entry</i> parameter is the entry retrieved from the directory. The <i>target</i> parameter should contain the attribute type the values that are required, including the optional language subtype. The <i>type</i> parameter points to a buffer that returns the attribute type retrieved by this function. Unlike the <code>ldap_get_values()</code> function, if a language subtype is specified, this function first attempts to find and return values that match that subtype, for example, <code>cn;lang-en</code>.</p> <p><code>ldap_get_lang_values_len()</code> returns a null-terminated array of pointers to <code>berval</code> structures, each containing the length and pointer to a binary value of an attribute for a given entry. The <i>entry</i> parameter is the result returned by <code>ldap_result()</code> or <code>ldap_search_s()</code> functions. The <i>target</i> parameter is the attribute returned by the call to <code>ldap_first_attribute()</code> or <code>ldap_next_attribute()</code>, or the attribute as a literal string, such as <code>jpegPhoto</code> or <code>audio</code>.</p> <p>These functions are deprecated. Use <code>ldap_get_values()</code> or <code>ldap_get_values_len()</code> instead.</p>
<b>RETURN VALUES</b>	<p>If successful, <code>ldap_get_lang_values()</code> returns a null-terminated array of the attribute’s values. If the call is unsuccessful, or if no such attribute exists in the <i>entry</i>, it returns a NULL and sets the appropriate error code in the LDAP structure.</p> <p>The <code>ldap_get_lang_values_len()</code> function returns a null-terminated array of pointers to <code>berval</code> structures, which in turn, if successful, contain pointers to the attribute’s binary values. If the call is unsuccessful, or if no such attribute exists in the <i>entry</i>, it returns a NULL and sets the appropriate error code in the LDAP structure.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ldap\_get\_lang\_values(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Obsolete

**SEE ALSO** ldap\_first\_attribute(3LDAP), ldap\_next\_attribute(3LDAP),  
ldap\_get\_values(3LDAP), ldap\_result(3LDAP), ldap\_search\_s(3LDAP),  
attributes(5)

<b>NAME</b>	ldap_get_option, ldap_set_option – get/set session preferences in the ldap structure.								
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  LDAP ldap_set_option(LDAP *ld, int option, void *optdata []); LDAP ldap_get_option(LDAP *ld, int option, void optdata []);</pre>								
<b>DESCRIPTION</b>	<p>These functions provide access to session preferences to an LDAP structure. ldap_get_option() gets session preferences from the LDAP structure. ldap_set_option() sets session preferences in the LDAP structure.</p> <p><i>ld</i> is the connection handle, which is a pointer to an LDAP structure containing information about the connection to the LDAP server. <i>option</i> is the name of the option to be read or modified. <i>optdata</i> is a pointer to the value of the option that you want to set/get.</p> <p>The <i>option</i> parameter can have one of the values listed in the following section.</p>								
<b>PARAMETERS</b>	<p>The following are the values for the <i>option</i> parameter:</p> <p><b>LDAP_OPT_API_INFO</b> Used to retrieve some basic information about the LDAP API implementation at execution time. The data type for the <i>optdata</i> parameter is (LDAPAPIInfo *). This option is READ-ONLY and cannot be set.</p> <p><b>LDAP_OPT_DEREF</b> Determines how aliases are handled during a search . The data type for the <i>optdata</i> parameter is (int *). <i>optdata</i> can be one of the following values:</p> <table border="0"> <tr> <td style="padding-right: 20px;">LDAP_DEREF_NEVER</td> <td>Specifies that aliases are never dereferenced.</td> </tr> <tr> <td style="padding-right: 20px;">LDAP_DEREF_SEARCHING</td> <td>Specifies that aliases are dereferenced when searching under the base object (but not when finding the base object).</td> </tr> <tr> <td style="padding-right: 20px;">LDAP_DEREF_FINDING</td> <td>Specifies that aliases are dereferenced when finding the base object (but not when searching under the base object).</td> </tr> <tr> <td style="padding-right: 20px;">LDAP_DEREF_ALWAYS</td> <td>Specifies that aliases are always dereferenced when finding the base object and searching under the base object.</td> </tr> </table> <p><b>LDAP_OPT_SIZELIMIT</b> Maximum number of entries that should be returned by the server in search results. The data type for the <i>optdata</i> parameter is (int *). Setting the <i>optdata</i> parameter to LDAP_NO_LIMIT removes any size limit enforced by the client.</p>	LDAP_DEREF_NEVER	Specifies that aliases are never dereferenced.	LDAP_DEREF_SEARCHING	Specifies that aliases are dereferenced when searching under the base object (but not when finding the base object).	LDAP_DEREF_FINDING	Specifies that aliases are dereferenced when finding the base object (but not when searching under the base object).	LDAP_DEREF_ALWAYS	Specifies that aliases are always dereferenced when finding the base object and searching under the base object.
LDAP_DEREF_NEVER	Specifies that aliases are never dereferenced.								
LDAP_DEREF_SEARCHING	Specifies that aliases are dereferenced when searching under the base object (but not when finding the base object).								
LDAP_DEREF_FINDING	Specifies that aliases are dereferenced when finding the base object (but not when searching under the base object).								
LDAP_DEREF_ALWAYS	Specifies that aliases are always dereferenced when finding the base object and searching under the base object.								

## ldap\_get\_option(3LDAP)

### LDAP\_OPT\_TIMELIMIT

Maximum number of seconds that should be spent by the server when answering a search request. The data type for the *optdata* parameter is (int \*). Setting the *optdata* parameter to LDAP\_NO\_LIMIT removes any time limit enforced by the client.

### LDAP\_OPT\_REFERRALS

Determines whether or not the client should follow referrals. The data type for the *optdata* parameter is (int \*). *optdata* can be one of the following values:

LDAP\_OPT\_ON                      Specifies that the client should follow referrals.

LDAP\_OPT\_OFF                     Specifies that the client should not follow referrals.

By default, the client follows referrals.

### LDAP\_OPT\_RESTART

Determines whether LDAP I/O operations are automatically restarted if they abort prematurely. It *may* be set to one of the constants LDAP\_OPT\_ON or LDAP\_OPT\_OFF.

### LDAP\_OPT\_PROTOCOL\_VERSION

Version of the protocol supported by your client. The data type for the *optdata* parameter is (int \*). You can specify either LDAP\_VERSION2 or LDAP\_VERSION3. If no version is set, the default is LDAP\_VERSION2. In order to use LDAP v3 features, you need to set the protocol version to LDAP\_VERSION3.

### LDAP\_OPT\_SERVER\_CONTROLS

Pointer to an array of LDAPControl structures representing the LDAP v3 server controls you want sent with every request by default. The data type for the *optdata* parameter for ldap\_set\_option() is (LDAPControl \*\*) and for ldap\_get\_option() is (LDAPControl \*\*\*).

### LDAP\_OPT\_CLIENT\_CONTROLS

Pointer to an array of LDAPControl structures representing the LDAP v3 client controls you want sent with every request by default. The data type for the *optdata* parameter for ldap\_set\_option() is (LDAPControl \*\*) and for ldap\_get\_option() is (LDAPControl \*\*\*).

### LDAP\_OPT\_API\_FEATURE\_INFO

Used to retrieve version information about LDAP API extended features at execution time. The data type for the *optdata* parameter is (LDAPAPIFeatureInfo \*). This option is READ-ONLY and cannot be set.

### LDAP\_OPT\_HOST\_NAME

This option sets the host name (or list of hosts) for the primary LDAP server. The data type for the *optdata* parameter for ldap\_set\_option() is (char \*), and for ldap\_get\_option() is (char \*\*).

### LDAP\_OPT\_ERROR\_NUMBER

The code of the most recent LDAP error that occurred for this session. The data type for the *optdata* parameter is (int \*).



## ldap\_get\_option(3LDAP)

### LDAP\_OPT\_ERROR\_STRING

The message returned with the most recent LDAP error that occurred for this session. The data type for the *optdata* parameter for `ldap_set_option()` is (char \*) and for `ldap_get_option()` is (char \*\*).

### LDAP\_OPT\_MATCHED\_DN

The matched DN value returned with the most recent LDAP error that occurred for this session. The data type for the *optdata* parameter for `ldap_set_option()` is (char \*) and for `ldap_get_option()` is (char \*\*).

### LDAP\_OPT\_REBIND\_ARG

Lets you set the last argument passed to the routine specified by `LDAP_OPT_REBIND_FN`. You can also set this option by calling the `ldap_set_rebind_proc()` function. The data type for the *optdata* parameter is (void \*).

### LDAP\_OPT\_REBIND\_FN

Lets you set the routine to be called when you need to authenticate a connection with another LDAP server (for example, during the course of a referral). You can also set this option by calling the `ldap_set_rebind_proc()` function. The data type for the *optdata* parameter is (LDAP\_REBINDPROC\_CALLBACK \*).

**RETURN VALUES** The `ldap_set_option()` and `ldap_get_option()` functions return:

<code>LDAP_SUCCESS</code>	If successful
<code>-1</code>	If unsuccessful

**ERRORS** Upon successful completion, both functions return `LDAP_SUCCESS`, otherwise `-1` is returned.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap_init(3LDAP)`, `attributes(5)`

**NOTES** There are other elements in the LDAP structure that you should not change. You should not make any assumptions about the order of elements in the LDAP structure.

## ldap\_get\_values(3LDAP)

<b>NAME</b>	ldap_get_values, ldap_get_values_len, ldap_count_values, ldap_count_values_len, ldap_value_free, ldap_value_free_len – LDAP attribute value handling functions				
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  char **ldap_get_values(LDAP *ld, LDAPMessage *entry, char *attr); struct berval **ldap_get_values_len(LDAP *ld, LDAPMessage *entry, char *attr);  ldap_count_values(char **vals); ldap_count_values_len(struct berval **vals); ldap_value_free(char **vals); ldap_value_free_len(struct berval **vals);</pre>				
<b>DESCRIPTION</b>	<p>These functions are used to retrieve and manipulate attribute values from an LDAP entry as returned by <code>ldap_first_entry(3LDAP)</code> or <code>ldap_next_entry(3LDAP)</code>. <code>ldap_get_values()</code> takes the <i>entry</i> and the attribute <i>attr</i> whose values are desired and returns a null-terminated array of the attribute's values. <i>attr</i> may be an attribute type as returned from <code>ldap_first_attribute(3LDAP)</code> or <code>ldap_next_attribute(3LDAP)</code>, or if the attribute type is known it can simply be given.</p> <p>The number of values in the array can be counted by calling <code>ldap_count_values()</code>. The array of values returned can be freed by calling <code>ldap_value_free()</code>.</p> <p>If the attribute values are binary in nature, and thus not suitable to be returned as an array of char *'s, the <code>ldap_get_values_len()</code> function can be used instead. It takes the same parameters as <code>ldap_get_values()</code>, but returns a null-terminated array of pointers to berval structures, each containing the length of and a pointer to a value.</p> <p>The number of values in the array can be counted by calling <code>ldap_count_values_len()</code>. The array of values returned can be freed by calling <code>ldap_value_free_len()</code>.</p>				
<b>ERRORS</b>	If an error occurs in <code>ldap_get_values()</code> or <code>ldap_get_values_len()</code> , NULL returned and the <code>ld_errno</code> field in the <code>ld</code> parameter is set to indicate the error. See <code>ldap_error(3LDAP)</code> for a description of possible error codes.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for a description of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWcsl (32-bit)</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsl (32-bit)
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsl (32-bit)				

ldap\_get\_values(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
	SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_first\_entry(3LDAP), ldap\_first\_attribute(3LDAP), ldap\_error(3LDAP), attributes(5)

**NOTES** These functions allocates memory that the caller must free.

## ldap\_memcache(3LDAP)

<b>NAME</b>	<code>ldap_memcache</code> , <code>ldap_memcache_init</code> , <code>ldap_memcache_set</code> , <code>ldap_memcache_get</code> , <code>ldap_memcache_flush</code> , <code>ldap_memcache_destroy</code> , <code>ldap_memcache_update</code> – LDAP client caching functions
<b>SYNOPSIS</b>	<pre>cc -flag ... file ...-lldap [-library ...] #include &lt;ldap.h&gt;  int ldap_memcache_init(unsigned long ttl, unsigned long size, char     **baseDNs, struct ldap_thread_fns *thread_fns, LDAPMemCache     **cachep);  int ldap_memcache_set(LDAP *ld, LDAPMemCache **cache);  int ldap_memcache_get(LDAP *ld, LDAPMemCache **cachep);  void ldap_memcache_flush(LDAPMemCache *cache, char *dn, int scope);  void ldap_memcache_destroy(LDAPMemCache *cache);  void ldap_memcache_update(LDAPMemCache *cache);</pre>
<b>DESCRIPTION</b>	<p>Use the <code>ldap_memcache</code> functions to maintain an in-memory client side cache to store search requests. Caching improves performance and reduces network bandwidth when a client makes repeated requests. The <i>cache</i> uses search criteria as the key to the cached items. When you send a search request, the <i>cache</i> checks the search criteria to determine if that request has been previously stored. If the request was stored, the search results are read from the <i>cache</i>.</p> <p>Make a call to <code>ldap_memcache_init()</code> to create the in-memory client side <i>cache</i>. The function passes back a pointer to an <code>LDAPMemCache</code> structure, which represents the <i>cache</i>. Make a call to the <code>ldap_memcache_set()</code> function to associate this <i>cache</i> with an LDAP connection handle, an LDAP structure. <i>ttl</i> is the the maximum amount of time (in seconds) that an item can be cached. If a <i>ttl</i> value of 0 is passed, there is no limit to the amount of time that an item can be cached. <i>size</i> is the maximum amount of memory (in bytes) that the cache will consume. A zero value of <i>size</i> means the cache has no size limit. <i>baseDNS</i> is an array of the base DN strings representing the base DN's of the search requests you want cached. If <i>baseDNS</i> is not <code>NULL</code>, only the search requests with the specified base DN's will be cached. If <i>baseDNS</i> is <code>NULL</code>, all search requests are cached. The <i>thread_fns</i> parameter takes an <code>ldap_thread_fns</code> structure specifying the functions that you want used to ensure that the cache is thread-safe. You should specify this if you have multiple threads that are using the same connection handle and cache. If you are not using multiple threads, pass <code>NULL</code> for this parameter.</p> <p><code>ldap_memcache_set()</code> associates an in-memory <i>cache</i> that you have already created by calling the <code>ldap_memcache_init()</code> function with an LDAP connection handle. The <i>ld</i> parameter should be the result of a successful call to <code>ldap_open(3LDAP)</code>. The <i>cache</i> parameter should be the result of a <i>cache</i> created by the <code>ldap_memcache_init()</code> call. After you call this function, search requests made over the specified LDAP connection will use this cache. To disassociate the cache from the LDAP connection handle, make a call to the <code>ldap_unbind(3LDAP)</code> or <code>ldap_unbind_ext(3LDAP)</code> function. Make a call to <code>ldap_memcache_set()</code> if you</p>

## ldap\_memcache(3LDAP)

want to associate a cache with multiple LDAP connection handles. For example, call the `ldap_memcache_get()` function to get the *cache* associated with one connection, then you can call this function and associate the *cache* with another connection.

The `ldap_memcache_get()` function gets the *cache* associated with the specified connection handle (LDAP structure). This *cache* is used by all search requests made through that connection. When you call this function, the function sets the *cachep* parameter as a pointer to the `LDAPMemCache` structure that is associated with the connection handle.

`ldap_memcache_flush()` flushes search requests from the *cache*. If the base DN of a search request is within the scope specified by the *dn* and *scope* arguments, the search request is flushed from the *cache*. If no DN is specified, the entire cache is flushed. The *scope* parameter, along with the *dn* parameter, identifies the search requests that you want flushed from the *cache*. This argument can have one of the following values:

`LDAP_SCOPE_BASE`  
`LDAP_SCOPE_ONELEVEL`  
`LDAP_SCOPE_SUBTREE`

`ldap_memcache_destroy()` frees the specified `LDAPMemCache` structure pointed to by *cache* from memory. Call this function after you are done working with a *cache*.

`ldap_memcache_update()` checks the cache for items that have expired and removes them. This check is typically done as part of the way the *cache* normally works. You do not need to call this function unless you want to update the *cache* at this point in time. This function is only useful in a multithreaded application, since it will not return until the *cache* is destroyed.

### PARAMETERS

<i>ttl</i>	The maximum amount of time (in seconds) that an item can be cached
<i>size</i>	The maximum amount of memory (in bytes) that the cache will consume.
<i>baseDNs</i>	An array of the base DN strings representing the base DN's of the search requests you want cached
<i>thread_fns</i>	A pointer to the <code>ldap_thread_fns</code> structure
<i>cachep</i>	A pointer to the <code>LDAPMemCache</code> structure
<i>cache</i>	The result of a <i>cache</i> created by the <code>ldap_memcache_init()</code> call
<i>ld</i>	The result of a successful call to <code>ldap_open(3LDAP)</code>
<i>dn</i>	The search requests that you want flushed from the <i>cache</i>

## ldap\_memcache(3LDAP)

*scope* The search requests that you want flushed from the *cache*

**ERRORS** The functions that have `int` return values return `LDAP_SUCCESS` if the operation was successful. Otherwise, they return another LDAP error code. See `ldap_error(3LDAP)` for a list of the LDAP error codes.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap_error(3LDAP)`, `ldap_open(3LDAP)`, `ldap_search(3LDAP)`, `attributes(5)`

<b>NAME</b>	ldap_memfree – free memory allocated by LDAP API functions								
<b>SYNOPSIS</b>	<pre>cc -flag ... file...-lldap [-library ...] #include &lt; lber.h&gt; #include &lt; ldap.h&gt;  void ldap_memfree(void *p);</pre>								
<b>DESCRIPTION</b>	<p>The ldap_memfree() function frees the memory allocated by certain LDAP API functions that do not have corresponding functions to free memory. These functions include ldap_get_dn(3LDAP), ldap_first_attribute(3LDAP), and ldap_next_attribute(3LDAP).</p> <p>The ldap_memfree() function takes one parameter, <i>p</i>, which is a pointer to the memory to be freed.</p>								
<b>PARAMETERS</b>	<i>p</i> A pointer to the memory to be freed.								
<b>RETURN VALUES</b>	There are no return values for the ldap_memfree() function.								
<b>ERRORS</b>	No errors are defined for the ldap_memfree() function.								
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:								
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWcsl (32-bit)</td> </tr> <tr> <td></td> <td>SUNWcslx (64-bit)</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsl (32-bit)		SUNWcslx (64-bit)	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWcsl (32-bit)								
	SUNWcslx (64-bit)								
Interface Stability	Evolving								
<b>SEE ALSO</b>	ldap(3ldap), ldap_first_attribute(3LDAP), ldap_get_dn(3LDAP), ldap_next_attribute(3LDAP), attributes(5)								

## ldap\_modify(3LDAP)

<b>NAME</b>	ldap_modify, ldap_modify_s, ldap_mods_free, ldap_modify_ext, ldap_modify_ext_s – LDAP entry modification functions
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_modify(LDAP *ld, char *dn, LDAPMod *mods[]); int ldap_modify_s(LDAP *ld, char *dn, LDAPMod *mods[]); void ldap_mods_free(LDAPMod **mods, int freemods); int ldap_modify_ext(LDAP *ld, char *dn, LDAPMod **mods, LDAPControl **serverctrls, LDAPControl **clientctrls, int *msgidp); int ldap_modify_ext_s(LDAP *ld, char *dn, LDAPMod **mods, LDAPControl **serverctrls, LDAPControl **clientctrls);</pre>
<b>DESCRIPTION</b>	<p>The function <code>ldap_modify_s()</code> is used to perform an LDAP modify operation. <i>dn</i> is the DN of the entry to modify, and <i>mods</i> is a null-terminated array of modifications to make to the entry. Each element of the <i>mods</i> array is a pointer to an LDAPMod structure, which is defined below.</p> <pre>typedef struct ldapmod {     int mod_op;     char *mod_type;     union {         char **modv_strvals;         struct berval **modv_bvals;     } mod_vals; } LDAPMod; #define mod_values mod_vals.modv_strvals #define mod_bvalues mod_vals.modv_bvals</pre> <p>The <i>mod_op</i> field is used to specify the type of modification to perform and should be one of LDAP_MOD_ADD, LDAP_MOD_DELETE, or LDAP_MOD_REPLACE. The <i>mod_type</i> and <i>mod_values</i> fields specify the attribute type to modify and a null-terminated array of values to add, delete, or replace respectively.</p> <p>If you need to specify a non-string value (for example, to add a photo or audio attribute value), you should set <i>mod_op</i> to the logical OR of the operation as above (for example, LDAP_MOD_REPLACE) and the constant LDAP_MOD_BVALUES. In this case, <i>mod_bvalues</i> should be used instead of <i>mod_values</i>, and it should point to a null-terminated array of struct berval, as defined in <code>&lt;lber.h&gt;</code>.</p> <p>For LDAP_MOD_ADD modifications, the given values are added to the entry, creating the attribute if necessary. For LDAP_MOD_DELETE modifications, the given values are deleted from the entry, removing the attribute if no values remain. If the entire attribute is to be deleted, the <i>mod_values</i> field should be set to NULL. For LDAP_MOD_REPLACE modifications, the attribute will have the listed values after the modification, having been created if necessary. All modifications are performed in the order in which they are listed.</p>



## ldap\_modify(3LDAP)

ldap\_modify\_s() returns the LDAP error code resulting from the modify operation.

The ldap\_modify() operation works the same way as ldap\_modify\_s(), except that it is asynchronous, returning the message id of the request it initiates, or -1 on error. The result of the operation can be obtained by calling ldap\_result(3LDAP).

ldap\_mods\_free() can be used to free each element of a null-terminated array of mod structures. If freemods is non-zero, the mods pointer itself is freed as well.

The ldap\_modify\_ext() function initiates an asynchronous modify operation and returns LDAP\_SUCCESS if the request was successfully sent to the server, or else it returns a LDAP error code if not. See ldap\_error(3LDAP). If successful, ldap\_modify\_ext() places the message id of the request in \*msgidp. A subsequent call to ldap\_result(3LDAP), can be used to obtain the result of the add request.

The ldap\_modify\_ext\_s() function initiates a synchronous modify operation and returns the result of the operation itself.

**ERRORS** ldap\_modify\_s() returns an LDAP error code, either LDAP\_SUCCESS or an error. See ldap\_error(3LDAP).

ldap\_modify() returns -1 in case of trouble, setting the error field of ld.

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_add(3LDAP), ldap\_error(3LDAP), ldap\_get\_option(3LDAP), attributes(5)

## ldap\_modrdn(3LDAP)

<b>NAME</b>	ldap_modrdn, ldap_modrdn_s, ldap_modrdn2, ldap_modrdn2_s, ldap_rename, ldap_rename_s – modify LDAP entry RDN
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_modrdn(LDAP *ld, const char *dn, const char *newrdn);  int ldap_modrdn_s(LDAP *ld, const char *dn, const char *newrdn, int   deleteoldrdn);  int ldap_modrdn2(LDAP *ld, const char *dn, const char *newrdn, int   deleteoldrdn);  int ldap_modrdn2_s(LDAP *ld, const char *dn, const char *newrdn, int   deleteoldrdn);  int ldap_rename(LDAP *ld, const char *dn, const char *newrdn, const   char *newparent, int deleteoldrdn, LDAPControl **serverctrls,   LDAPControl **clientctrls, int *msgidp);  int ldap_rename_s(LDAP *ld, const char *dn, const char *newrdn,   const char *newparent, const int deleteoldrdn, LDAPControl   **serverctrls, LDAPControl **clientctrls);</pre>
<b>DESCRIPTION</b>	<p>The <code>ldap_modrdn()</code> and <code>ldap_modrdn_s()</code> functions perform an LDAP modify RDN (Relative Distinguished Name) operation. They both take <code>dn</code>, the DN (Distinguished Name) of the entry whose RDN is to be changed, and <code>newrdn</code>, the new RDN, to give the entry. The old RDN of the entry is never kept as an attribute of the entry. <code>ldap_modrdn()</code> is asynchronous. It returns the message id of the operation it initiates. <code>ldap_modrdn_s()</code> is synchronous. It returns the LDAP error code that indicates the success or failure of the operation.</p> <p>The <code>ldap_modrdn2()</code> and <code>ldap_modrdn2_s()</code> functions also perform an LDAP modify RDN operation. They take the same parameters as above. In addition, they both take the <code>deleteoldrdn</code> parameter, which is used as a boolean value to indicate whether or not the old RDN values should be deleted from the entry.</p> <p>The <code>ldap_rename()</code>, <code>ldap_rename_s()</code> routines are used to change the name, that is, the RDN of an entry. These routines deprecate the <code>ldap_modrdn()</code> and <code>ldap_modrdn_s()</code> routines, as well as <code>ldap_modrdn2()</code> and <code>ldap_modrdn2_s()</code>.</p> <p>The <code>ldap_rename()</code> and <code>ldap_rename_s()</code> functions both support LDAPv3 server controls and client controls.</p>
<b>ERRORS</b>	<p>The synchronous (<code>_s</code>) versions of these functions return an LDAP error code, either <code>LDAP_SUCCESS</code> or an error. See <code>ldap_error(3LDAP)</code>.</p> <p>The asynchronous versions return <code>-1</code> in the event of an error, setting the <code>ld_errno</code> field of <code>ld</code>. See <code>ldap_error(3LDAP)</code> for more details. Use <code>ldap_result(3LDAP)</code> to determine a particular unsuccessful result.</p>

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes of the `ldap_modrdn()`, `ldap_modrdn_s()`, `ldap_modrdn2()` and `ldap_modrdn2_s()` functions:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Obsolete

The `ldap_rename()` and `ldap_rename_s()` functions have the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap(3LDAP)`, `ldap_error(3LDAP)`, `attributes(5)`

## ldap\_open(3LDAP)

<b>NAME</b>	ldap_open, ldap_init – initialize an LDAP session				
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  LDAP *ldap_open(const char *host, int port); LDAP *ldap_init(const char *host, int port);</pre>				
<b>DESCRIPTION</b>	<p>ldap_open() initializes an LDAP session and additionally opens a connection to an LDAP server before it returns to the caller. Unlike ldap_open(), ldap_init() does not open a connection to the LDAP server until an operation is performed that requires it, for example, a search request.</p> <p>The ldap_open() function is deprecated and should no longer be used. Call ldap_init() instead.</p> <p>Either function takes a space-separated list of hostnames or strings in the dot-notation form that represents the IP addresses of hosts running LDAP servers to which it can connect. Each of these hostnames may include a port number, which is separated from the hostname with a colon (":") character. If the port number is specified as part of the hostname, then this takes precedence over the port parameter. The functions will try these hosts in the order listed, and return the first successful connection.</p>				
<b>PARAMETERS</b>	<p>These functions support the following parameters.</p> <table><tr><td><i>host</i></td><td>A space-separated list of hostnames or strings in the dot-notation form that represents the IP addresses of hosts running LDAP servers to which it can connect.</td></tr><tr><td><i>port</i></td><td>The TCP port number to which to connect. Supply the constant LDAP_PORT to obtain the default LDAP port of 389. If a host includes a port number, then this parameter is ignored.</td></tr></table>	<i>host</i>	A space-separated list of hostnames or strings in the dot-notation form that represents the IP addresses of hosts running LDAP servers to which it can connect.	<i>port</i>	The TCP port number to which to connect. Supply the constant LDAP_PORT to obtain the default LDAP port of 389. If a host includes a port number, then this parameter is ignored.
<i>host</i>	A space-separated list of hostnames or strings in the dot-notation form that represents the IP addresses of hosts running LDAP servers to which it can connect.				
<i>port</i>	The TCP port number to which to connect. Supply the constant LDAP_PORT to obtain the default LDAP port of 389. If a host includes a port number, then this parameter is ignored.				
<b>ERRORS</b>	<p>Either function returns a handle to an LDAP session that contains a pointer to an opaque structure that must be passed to subsequent calls for that session. If these sessions cannot be initialized, these functions will return NULL, and errno should be set appropriately.</p> <p>Various aspects of this opaque structure can be read or written to control the session-wide parameters. Use ldap_get_option(3LDAP) to access the current value of these options and ldap_set_option(3LDAP) to set the value of these options.</p>				
<b>ATTRIBUTES</b>	See attributes(5) for a description of the following attributes:				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)

ldap\_open(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
	SUNWcslx (64-bit)
Stability Level	Evolving

**SEE ALSO** `errno(3C)`, `ldap(3LDAP)`, `ldap_bind(3LDAP)`, `ldap_get_option(3LDAP)`, `ldap_set_option(3LDAP)`, `attributes(5)`

## ldap\_parse\_result(3LDAP)

<b>NAME</b>	ldap_parse_result, ldap_parse_extended_result, ldap_parse_sasl_bind_result – LDAP message result parser						
<b>SYNOPSIS</b>	<pre>cc[ <i>flag...</i> ] <i>file...</i> -lldap[ <i>library...</i> ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_parse_result(LDAP *ld, LDAPMessage *res, int *errcodep, char     **matcheddn, char **errmsgp, char ***referralsp, LDAPControl     ***serverctrlsp, int freeit);  int ldap_parse_sasl_bind_result(LDAP *ld, LDAPMessage *res, struct     berval**servercredp, int freeit);  int ldap_parse_extended_result(LDAP *ld, LDAPMessage *res, char     **resultoidp, struct berval **resultdata, int freeit);</pre>						
<b>DESCRIPTION</b>	The ldap_parse_extended_result(), ldap_parse_result() and ldap_parse_sasl_bind_result() routines search for a message to parse. These functions skip messages of type LDAP_RES_SEARCH_ENTRY and LDAP_RES_SEARCH_REFERENCE.						
<b>RETURN VALUES</b>	They return LDAP_SUCCESS if the result was successfully parsed or an LDAP error code if not (see ldap_error(3LDAP)).						
<b>ATTRIBUTES</b>	See attributes(5) for a description of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWcsl (32-bit) SUNWcslx (64-bit)</td></tr><tr><td>Interface Stability</td><td>Evolving</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)						
Interface Stability	Evolving						
<b>SEE ALSO</b>	ldap_error(3LDAP), ldap_result(3LDAP), attributes(5)						

<b>NAME</b>	ldap_result, ldap_msgfree – wait for and return LDAP operation result
<b>SYNOPSIS</b>	<pre>cc[ flag... ] file... -lldap[ library... ] #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_result(LDAP *ld, int msgid, int all, struct timeval *timeout,                LDAPMessage **result);  int ldap_msgfree(LDAPMessage *msg);</pre>
<b>DESCRIPTION</b>	<p>The <code>ldap_result()</code> function is used to wait for and return the result of an operation previously initiated by one of the LDAP asynchronous operation functions, for example, <code>ldap_search(3LDAP)</code>, and <code>ldap_modify(3LDAP)</code>. Those functions all return <code>-1</code> in case of error, and an invocation identifier upon successful initiation of the operation. The invocation identifier is picked by the library and is guaranteed to be unique across the LDAP session. It can be used to request the result of a specific operation from <code>ldap_result()</code> through the <code>msgid</code> parameter.</p> <p>The <code>ldap_result()</code> function will block or not, depending upon the setting of the <code>timeout</code> parameter. If <code>timeout</code> is not a null pointer, it specifies a maximum interval to wait for the selection to complete. If <code>timeout</code> is a null pointer, the select blocks indefinitely. To effect a poll, the <code>timeout</code> argument should be a non-null pointer, pointing to a zero-valued <code>timeval</code> structure. See <code>select(1)</code> for further details.</p> <p>If the result of a specific operation is required, <code>msgid</code> should be set to the invocation identifier returned when the operation was initiated, otherwise <code>LDAP_RES_ANY</code> should be supplied. The <code>all</code> parameter only has meaning for search responses and is used to select whether a single entry of the search response should be returned, or all results of the search should be returned.</p> <p>A search response is made up of zero or more search entries followed by a search result. If <code>all</code> is set to <code>-</code>, search entries will be returned one at a time as they come in, by means of separate calls to <code>ldap_result()</code>. If it is set to <code>-1</code>, the search response will only be returned in its entirety, that is, after all entries and the final search result have been received.</p> <p>Upon success, the type of the result received is returned and the <code>result</code> parameter will contain the result of the operation. This result should be passed to the LDAP parsing functions, (see <code>ldap_first_entry(3LDAP)</code>) for interpretation.</p> <p>The possible result types returned are:</p> <pre>#define LDAP_RES_BIND          0x61L #define LDAP_RES_SEARCH_ENTRY  0x64L #define LDAP_RES_SEARCH_RESULT 0x65L #define LDAP_RES_MODIFY        0x67L #define LDAP_RES_ADD           0x69L #define LDAP_RES_DELETE        0x6bL #define LDAP_RES_MODRDN        0x6dL #define LDAP_RES_COMPARE       0x6fL</pre>

## ldap\_result(3LDAP)

The `ldap_msgfree()` function is used to free the memory allocated for a result by `ldap_result()` or `ldap_search_s(3LDAP)` functions. It takes a pointer to the result to be freed and returns the type of the message it freed.

**ERRORS** `ldap_result()` returns `-1` if something bad happens, and zero if the timeout specified was exceeded.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `select(1)`, `ldap(3LDAP)`, `ldap_search(3LDAP)`, `attributes(5)`

**NOTES** This function allocates memory for results that it receives. The memory can be freed by calling `ldap_msgfree`.



<b>NAME</b>	ldap_search, ldap_search_s, ldap_search_ext, ldap_search_ext_s, ldap_search_st – LDAP search operations
<b>SYNOPSIS</b>	<pre>cc[ flag... ] file... -lldap[ library... ]  #include &lt;sys/time.h&gt; /* for struct timeval definition */ #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_search(LDAP *ld, char *base, int scope, char *filter, char                *attrs[], int attrsonly);  int ldap_search_s(LDAP *ld, char *base, int scope, char *filter, char                  *attrs[], int attrsonly, LDAPMessage **res);  int ldap_search_st(LDAP *ld, char *base, int scope, char *filter, char                   *attrs[], int attrsonly, struct timeval *timeout, LDAPMessage **res);  int ldap_search_ext(LDAP *ld, char *base, int scope, char *filter, char                    **attrs, int attrsonly, LDAPControl **serverctrls, LDAPControl                    **clientctrls, struct timeval *timeoutp, int sizelimit, int *msgidp);  int ldap_search_ext_s(LDAP *ld, char *base, int scope, char *filter,                      char **attrs, int attrsonly, LDAPControl **serverctrls, LDAPControl                      **clientctrls, struct timeval *timeoutp, int sizelimit);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to perform LDAP search operations. <code>ldap_search_s()</code> does the search synchronously (that is, not returning until the operation completes). <code>ldap_search_st()</code> does the same, but allows a <i>timeout</i> to be specified. <code>ldap_search()</code> is the asynchronous version, initiating the search and returning the message id of the operation it initiated.</p> <p><i>Base</i> is the DN of the entry at which to start the search. <i>Scope</i> is the scope of the search and should be one of <code>LDAP_SCOPE_BASE</code>, to search the object itself, <code>LDAP_SCOPE_ONELEVEL</code>, to search the object's immediate children, or <code>LDAP_SCOPE_SUBTREE</code>, to search the object and all its descendants.</p> <p><i>Filter</i> is a string representation of the filter to apply in the search. Simple filters can be specified as <i>attributetype=attributevalue</i>. More complex filters are specified using a prefix notation according to the following BNF:</p> <pre>&lt;filter&gt; ::= '(' &lt;filtercomp&gt; ')' &lt;filtercomp&gt; ::= &lt;and&gt;   &lt;or&gt;   &lt;not&gt;   &lt;simple&gt; &lt;and&gt; ::= '&amp;' &lt;filterlist&gt; &lt;or&gt; ::= ' ' &lt;filterlist&gt; &lt;not&gt; ::= '!' &lt;filter&gt; &lt;filterlist&gt; ::= &lt;filter&gt;   &lt;filter&gt; &lt;filterlist&gt; &lt;simple&gt; ::= &lt;attributetype&gt; &lt;filtertype&gt; &lt;attributevalue&gt; &lt;filtertype&gt; ::= '='   '~='   '&lt;'   '&gt;'</pre> <p>The <code>'~='</code> construct is used to specify approximate matching. The representation for <code>&lt;attributetype&gt;</code> and <code>&lt;attributevalue&gt;</code> are as described in RFC 1778. In addition, <code>&lt;attributevalue&gt;</code> can be a single <code>*</code> to achieve an attribute existence test, or can contain text and <code>*</code>'s interspersed to achieve substring matching.</p>

## ldap\_search(3LDAP)

For example, the filter "mail=\*" will find any entries that have a mail attribute. The filter "mail=\*@terminator.rs.itd.umich.edu" will find any entries that have a mail attribute ending in the specified string. To put parentheses in a filter, escape them with a backslash '\ ' character. See RFC 1588 for a more complete description of allowable filters. See `ldap_getfilter(3LDAP)` for functions to help in constructing search filters automatically.

*Attrs* is a null-terminated array of attribute types to return from entries that match *filter*. If NULL is specified, all attributes will be returned. *Attronly* should be set to 1 if only attribute types are wanted. It should be set to 0 if both attributes types and attribute values are wanted.

The `ldap_search_ext()` function initiates an asynchronous search operation and returns LDAP\_SUCCESS if the request was successfully sent to the server, or else it returns a LDAP error code (see `ldap_error(3LDAP)`). If successful, `ldap_search_ext()` places the message id of the request in *\*msgidp*. A subsequent call to `ldap_result(3LDAP)`, can be used to obtain the result of the add request.

The `ldap_search_ext_s()` function initiates a synchronous search operation and as such returns the result of the operation itself.

**ERRORS** `ldap_search_s()` and `ldap_search_st()` will return the LDAP error code resulting from the search operation. See `ldap_error(3LDAP)` for details.

`ldap_search()` returns -1 when terminating unsuccessfully.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `ldap(3LDAP)`, `ldap_result(3LDAP)`, `ldap_getfilter(3LDAP)`, `ldap_error(3LDAP)`, `attributes(5)`

**NOTES** Note that both read and list functionality are subsumed by these functions, by using a filter like "objectclass=\*" and a scope of LDAP\_SCOPE\_BASE (to emulate read) or LDAP\_SCOPE\_ONELEVEL (to emulate list).

These functions may allocate memory which must be freed by the calling application. Return values are contained in `<ldap.h>`.

<b>NAME</b>	ldap_searchprefs, ldap_init_searchprefs, ldap_init_searchprefs_buf, ldap_free_searchprefs, ldap_first_searchobj, ldap_next_searchobj – LDAP search preference configuration routines
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  # include &lt;lber.h&gt; # include &lt;ldap.h&gt;  int ldap_init_searchprefs(char **file, struct ldap_searchobj     **solist);  int ldap_init_searchprefs_buf(char **buf, unsigned longlen, struct     ldap_searchobj **solist);  struct ldap_searchobj **ldap_free_searchprefs(struct     ldap_searchobj **solist);  struct ldap_searchobj **ldap_first_searchobj(struct     ldap_seachobj **solist);  struct ldap_searchobj **ldap_next_searchobj(struct ldap_seachobj     **solist, struct ldap_seachobj **so);</pre>
<b>DESCRIPTION</b>	<p>These functions provide a standard way to access LDAP search preference configuration data. LDAP search preference configurations are typically used by LDAP client programs to specify which attributes a user may search by, labels for the attributes, and LDAP filters and scopes associated with those searches. Client software presents these choices to a user, who can then specify the type of search to be performed.</p> <p>ldap_init_searchprefs() reads a sequence of search preference configurations from a valid LDAP searchpref configuration file. See ldapsearchprefs.conf(4). Upon success, 0 is returned and <i>solist</i> is set to point to a list of search preference data structures.</p> <p>ldap_init_searchprefs_buf() reads a sequence of search preference configurations from <i>buf</i>, whose size is <i>buflen</i>. <i>buf</i> should point to the data in the format defined for an LDAP search preference configuration file. See ldapsearchprefs.conf(4). Upon success, 0 is returned and <i>solist</i> is set to point to a list of search preference data structures.</p> <p>ldap_free_searchprefs() disposes of the data structures allocated by ldap_init_searchprefs().</p> <p>ldap_first_searchpref() returns the first search preference data structure in the list <i>solist</i>. The <i>solist</i> is typically obtained by calling ldap_init_searchprefs().</p> <p>ldap_next_searchpref() returns the search preference after <i>so</i> in the template list <i>solist</i>. A NULL pointer is returned if <i>so</i> is the last entry in the list.</p>
<b>ERRORS</b>	ldap_init_search_prefs() and ldap_init_search_prefs_bufs() return:

## ldap\_searchprefs(3LDAP)

LDAP\_SEARCHPREF\_ERR\_VERSION      *\*\*buf* points to data that is newer than can be handled.

LDAP\_SEARCHPREF\_ERR\_MEM          Memory allocation problem.

**ATTRIBUTES**      See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO**      `ldap(3LDAP)`, `ldapsearchprefs.conf(4)`, `attributes(5)`

Yeong, W., Howes, T., and Hardcastle-Kille, S., "Lightweight Directory Access Protocol", OSI-DS-26, April 1992.

Howes, T., Hardcastle-Kille, S., Yeong, W., and Robbins, C., "Lightweight Directory Access Protocol", OSI-DS-26, April 1992.

Hardcastle-Kille, S., "A String Representation of Distinguished Names", OSI-DS-23, April 1992.

Information Processing - Open Systems Interconnection - The Directory, International Organization for Standardization. International Standard 9594, (1988).

<b>NAME</b>	ldap_sort, ldap_sort_entries, ldap_sort_values, ldap_sort_strcasecmp – LDAP entry sorting functions
<b>SYNOPSIS</b>	<pre>cc[ <i>flag...</i> ] <i>file...</i> -lldap[ <i>library...</i> ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  <b>ldap_sort_entries</b>(LDAP *<i>ld</i>, LDAPMessage **<i>chain</i>, char *<i>attr</i>, int     (*<i>cmp</i>) ());  <b>ldap_sort_values</b>(LDAP *<i>ld</i>, char **<i>vals</i>, int (*<i>cmp</i>) ());  <b>ldap_sort_strcasecmp</b>(char *<i>a</i>, char *<i>b</i>);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to sort lists of entries and values retrieved from an LDAP server. <code>ldap_sort_entries()</code> is used to sort a chain of entries retrieved from an LDAP search call either by DN or by some arbitrary attribute in the entries. It takes <i>ld</i>, the LDAP structure, which is only used for error reporting, <i>chain</i>, the list of entries as returned by <code>ldap_search_s(3LDAP)</code> or <code>ldap_result(3LDAP)</code>. <i>attr</i> is the attribute to use as a key in the sort or NULL to sort by DN, and <i>cmp</i> is the comparison function to use when comparing values (or individual DN components if sorting by DN). In this case, <i>cmp</i> should be a function taking two single values of the <i>attr</i> to sort by, and returning a value less than zero, equal to zero, or greater than zero, depending on whether the first argument is less than, equal to, or greater than the second argument. The convention is the same as used by <code>qsort(3C)</code>, which is called to do the actual sorting.</p> <p><code>ldap_sort_values()</code> is used to sort an array of values from an entry, as returned by <code>ldap_get_values(3LDAP)</code>. It takes the LDAP connection structure <i>ld</i>, the array of values to sort <i>vals</i>, and <i>cmp</i>, the comparison function to use during the sort. Note that <i>cmp</i> will be passed a pointer to each element in the <i>vals</i> array, so if you pass the normal <code>char **</code> for this parameter, <i>cmp</i> should take two <code>char **</code>'s as arguments (that is, you cannot pass <code>strcasecmp</code> or its friends for <i>cmp</i>). You can, however, pass the function <code>ldap_sort_strcasecmp()</code> for this purpose.</p> <p>For example:</p> <pre>LDAP *ld; LDAPMessage *res; /* ... call to ldap_search_s(), fill in res, retrieve sn attr ... */  /* now sort the entries on surname attribute */ if ( ldap_sort_entries( ld, &amp;res, "sn", ldap_sort_strcasecmp ) != 0 )     ldap_perror( ld, "ldap_sort_entries" );</pre>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)

## ldap\_sort(3LDAP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
	SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_search(3LDAP), ldap\_result(3LDAP), qsort(3C), attributes(5)

**NOTES** The `ldap_sort_entries()` function applies the comparison function to each value of the attribute in the array as returned by a call to `ldap_get_values(3LDAP)`, until a mismatch is found. This works fine for single-valued attributes, but may produce unexpected results for multi-valued attributes. When sorting by DN, the comparison function is applied to an exploded version of the DN, without types. The return values for all of these functions are declared in the `<ldap.h>` header file. Some functions may allocate memory which must be freed by the calling application.

<b>NAME</b>	ldap_ufn, ldap_ufn_search_s, ldap_ufn_search_c, ldap_ufn_search_ct, ldap_ufn_setfilter, ldap_ufn_setprefix, ldap_ufn_timeout – LDAP user friendly search functions
<b>SYNOPSIS</b>	<pre>cc[ flag... ] file... -lldap[ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_ufn_search_c(LDAP *ld, char *ufn, char **attrs, int attrsonly,     LDAPMessage **res, int (*cancelproc)(), void *cancelparm);  int ldap_ufn_search_ct(LDAP *ld, char *ufn, char **attrs, int attrsonly,     LDAPMessage **res, int (*cancelproc)(), void *cancelparm, char *tag1,     char *tag2, char *tag3);  int ldap_ufn_search_s(LDAP *ld, char *ufn, char **attrs, int attrsonly,     LDAPMessage **res);  LDAPFilterDesc *ldap_ufn_setfilter(LDAP *ld, char *fname);  void ldap_ufn_setprefix(LDAP *ld, char *prefix);  int ldap_ufn_timeout(void *toparam);</pre>
<b>DESCRIPTION</b>	<p>These functions are used to perform LDAP user friendly search operations. ldap_ufn_search_s() is the simplest form. It does the search synchronously. It takes ld to identify the the LDAP connection. The ufn parameter is the user friendly name for which to search. The attrs, attrsonly and res parameters are the same as for ldap_search(3LDAP).</p> <p>The ldap_ufn_search_c() function functions the same as ldap_ufn_search_s(), except that it takes cancelproc, a function to call periodically during the search. It should be a function taking a single void * argument, given by cancelparm. If cancelproc returns a non-zero result, the search will be abandoned and no results returned. The purpose of this function is to provide a way for the search to be cancelled, for example, by a user or because some other condition occurs.</p> <p>The ldap_ufn_search_ct() function is like ldap_ufn_search_c(), except that it takes three extra parameters. tag1 is passed to the ldap_init_getfilter(3LDAP) function when resolving the first component of the UFN. tag2 is used when resolving intermediate components. tag3 is used when resolving the last component. By default, the tags used by the other UFN search functions during these three phases of the search are "ufn first", "ufn intermediate", and "ufn last".</p> <p>The ldap_ufn_setfilter() function is used to set the ldapfilter.conf(4) file for use with the ldap_init_getfilter(3LDAP) function to fname.</p>

## ldap\_ufn(3LDAP)

The `ldap_ufn_setprefix()` function is used to set the default prefix (actually, it's a suffix) appended to UFNs before searching. UFNs with fewer than three components have the prefix appended first, before searching. If that fails, the UFN is tried with progressively shorter versions of the prefix, stripping off components. If the UFN has three or more components, it is tried by itself first. If that fails, a similar process is applied with the prefix appended.

The `ldap_ufn_timeout()` function is used to set the timeout associated with `ldap_ufn_search_s()` searches. The *timeout* parameter should actually be a pointer to a struct `timeval`. This is so `ldap_ufn_timeout()` can be used as a cancelproc in the above functions.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** `gettimeofday(3C)`, `ldap(3LDAP)`, `ldap_search(3LDAP)`, `ldap_getfilter(3LDAP)`, `ldapfilter.conf(4)`, `ldap_error(3LDAP)`, `attributes(5)`

**NOTES** These functions may allocate memory. Return values are contained in `<ldap.h>`.



<b>NAME</b>	ldap_url, ldap_is_ldap_url, ldap_url_parse, ldap_free_urldesc, ldap_url_search, ldap_url_search_s, ldap_url_search_st, ldap_dns_to_url, ldap_dn_to_url – LDAP Uniform Resource Locator functions
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lldap [ library... ]  #include &lt;lber.h&gt; #include &lt;ldap.h&gt;  int ldap_is_ldap_url(char *url);  int ldap_url_parse(char *url, LDAPURLDesc **ludpp);  ldap_free_urldesc(LDAPURLDesc *ludp);  int ldap_url_search(LDAP *ld, char *url, int attrsonly);  int ldap_url_search_s(LDAP *ld, char *url, int attrsonly, LDAPMessage **res);  int ldap_url_search_st(LDAP *ld, char *url, int attrsonly, struct     timeval *timeout, LDAPMessage **res);  char *ldap_dns_to_url(LDAP *ld, char *dns_name, char *attrs, char     *scope, char *filter);  char *ldap_dn_to_url(LDAP *ld, char *dn, int nameparts);</pre>
<b>DESCRIPTION</b>	<p>These functions support the use of LDAP URLs (Uniform Resource Locators). LDAP URLs look like this:</p> <pre>ldap://hostport/dn[?attributes[?scope[?filter]]]</pre> <p>where:</p> <p><i>hostport</i>            Host name with an optional ":portnumber".</p> <p><i>dn</i>                    Base DN to be used for an LDAP search operation.</p> <p><i>attributes</i>           Comma separated list of attributes to be retrieved.</p> <p><i>scope</i>                One of these three strings: base one sub (default=base).</p> <p><i>filter</i>                LDAP search filter as used in a call to ldap_search(3LDAP).</p> <p>Here is an example:</p> <pre>ldap://ldap.itd.umich.edu/c=US?o,description?one?o=umich</pre> <p>URLs that are wrapped in angle-brackets and/or preceded by "URL:" are also tolerated.</p> <p>ldap_is_ldap_url() returns a non-zero value if <i>url</i> looks like an LDAP URL (as opposed to some other kind of URL). It can be used as a quick check for an LDAP URL; the ldap_url_parse() function should be used if a more thorough check is needed.</p>

## ldap\_url(3LDAP)

`ldap_url_parse()` breaks down an LDAP URL passed in *url* into its component pieces. If successful, zero is returned, an LDAP URL description is allocated, filled in, and *ludpp* is set to point to it. See RETURN VALUES for values returned upon error.

`ldap_free_urldesc()` should be called to free an LDAP URL description that was obtained from a call to `ldap_url_parse()`.

`ldap_url_search()` initiates an asynchronous LDAP search based on the contents of the *url* string. This function acts just like `ldap_search(3LDAP)` except that many search parameters are pulled out of the URL.

`ldap_url_search_s()` performs a synchronous LDAP search based on the contents of the *url* string. This function acts just like `ldap_search_s(3LDAP)` except that many search parameters are pulled out of the URL.

`ldap_url_search_st()` performs a synchronous LDAP URL search with a specified *timeout*. This function acts just like `ldap_search_st(3LDAP)` except that many search parameters are pulled out of the URL.

`ldap_dns_to_url()` locates the LDAP URL associated with a DNS domain name. The supplied DNS domain name is converted into a distinguished name. The directory entry specified by that distinguished name is searched for a labeledURI attribute. If successful then the corresponding LDAP URL is returned. If unsuccessful then that entry's parent is searched and so on until the target distinguished name is reduced to only two nameparts. If *dns\_name* is NULL then the environment variable LOCALDOMAIN is used. If *attrs* is not NULL then it is appended to the URL's attribute list. If *scope* is not NULL then it overrides the URL's scope. If *filter* is not NULL then it is merged with the URL's filter. If an error is encountered then zero is returned, otherwise a string URL is returned. The caller should free the returned string if it is non-zero.

`ldap_dn_to_url()` locates the LDAP URL associated with a distinguished name. The number of nameparts in the supplied distinguished name must be provided. The specified directory entry is searched for a labeledURI attribute. If successful then the LDAP URL is returned. If unsuccessful then that entry's parent is searched and so on until the target distinguished name is reduced to only two nameparts. If an error is encountered then zero is returned, otherwise a string URL is returned. The caller should free the returned string if it is non-zero.

### RETURN VALUES

Upon error, one of these values is returned for `ldap_url_parse()`:

LDAP_URL_ERR_NOTLDAP	URL doesn't begin with "ldap://".
LDAP_URL_ERR_NODN	URL has no DN (required).
LDAP_URL_ERR_BADSCOPE	URL scope string is invalid.
LDAP_URL_ERR_MEM	Can't allocate memory space.

### ATTRIBUTES

See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving

**SEE ALSO** ldap(3LDAP), ldap\_search(3LDAP), attributes(5)

An LDAP URL Format , Tim Howes and Mark Smith, December 1995. Internet Draft (work in progress). Currently available at this URL:

<ftp://ds.internic.net/internet-drafts/draft-ietf-asid-ldap-format-03.txt>

## ldap\_version(3LDAP)

**NAME** ldap\_version – get version information about the LDAP SDK for C

**SYNOPSIS**

```
cc -flag ... file ...-lldap [-library ...]
#include <ldap.h>

int ldap_version(LDAPVERSION *ver);
```

**DESCRIPTION** A call to this function returns the version information for the LDAP SDK for C. This is a deprecated function. Use ldap\_get\_option(LDAP) instead. The version information is returned in the LDAPVersion structure pointed to by *ver*. If NULL is passed for *ver*, then only the SDK version will be returned.

**RETURN VALUES** The ldap\_version() function returns the version number of the LDAP SDK for C, multiplied by 100. For example, for version 1.0 of the LDAP SDK for C, the function returns 100.

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Obsolete

**SEE ALSO** ldap\_get\_option(LDAP), attributes(5)

<b>NAME</b>	listen – listen for connections on a socket						
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  int listen(int s, int backlog);</pre>						
<b>DESCRIPTION</b>	<p>To accept connections, a socket is first created with <code>socket(3SOCKET)</code>, a backlog for incoming connections is specified with <code>listen()</code> and then the connections are accepted with <code>accept(3SOCKET)</code>. The <code>listen()</code> call applies only to sockets of type <code>SOCK_STREAM</code> or <code>SOCK_SEQPACKET</code>.</p> <p>The <i>backlog</i> parameter defines the maximum length the queue of pending connections may grow to.</p> <p>If a connection request arrives with the queue full, the client will receive an error with an indication of <code>ECONNREFUSED</code> for <code>AF_UNIX</code> sockets. If the underlying protocol supports retransmission, the connection request may be ignored so that retries may succeed. For <code>AF_INET</code> and <code>AF_INET6</code> sockets, the TCP will retry the connection. If the <i>backlog</i> is not cleared by the time the tcp times out, the connect will fail with <code>ETIMEDOUT</code>.</p>						
<b>RETURN VALUES</b>	A 0 return value indicates success; -1 indicates an error.						
<b>ERRORS</b>	<p>The call fails if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>EBADF</code></td> <td>The argument <i>s</i> is not a valid file descriptor.</td> </tr> <tr> <td><code>ENOTSOCK</code></td> <td>The argument <i>s</i> is not a socket.</td> </tr> <tr> <td><code>EOPNOTSUPP</code></td> <td>The socket is not of a type that supports the operation <code>listen()</code>.</td> </tr> </table>	<code>EBADF</code>	The argument <i>s</i> is not a valid file descriptor.	<code>ENOTSOCK</code>	The argument <i>s</i> is not a socket.	<code>EOPNOTSUPP</code>	The socket is not of a type that supports the operation <code>listen()</code> .
<code>EBADF</code>	The argument <i>s</i> is not a valid file descriptor.						
<code>ENOTSOCK</code>	The argument <i>s</i> is not a socket.						
<code>EOPNOTSUPP</code>	The socket is not of a type that supports the operation <code>listen()</code> .						
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
MT-Level	Safe						
<b>SEE ALSO</b>	<code>accept(3SOCKET)</code> , <code>connect(3SOCKET)</code> , <code>socket(3SOCKET)</code> , <code>attributes(5)</code> , <code>socket(3HEAD)</code>						
<b>NOTES</b>	There is currently no <i>backlog</i> limit.						

## listen(3XNET)

<b>NAME</b>	listen – listen for socket connections and limit the queue of incoming connections																
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;sys/socket.h&gt;  int listen(int socket, int backlog);</pre>																
<b>DESCRIPTION</b>	<p>The <code>listen()</code> function marks a connection-mode socket, specified by the <i>socket</i> argument, as accepting connections, and limits the number of outstanding connections in the socket's listen queue to the value specified by the <i>backlog</i> argument.</p> <p>If <code>listen()</code> is called with a <i>backlog</i> argument value that is less than 0, the function sets the length of the socket's listen queue to 0.</p> <p>The implementation may include incomplete connections in the queue subject to the queue limit. The implementation may also increase the specified queue limit internally if it includes such incomplete connections in the queue subject to this limit.</p> <p>Implementations may limit the length of the socket's listen queue. If <i>backlog</i> exceeds the implementation-dependent maximum queue length, the length of the socket's listen queue will be set to the maximum supported value.</p> <p>The socket in use may require the process to have appropriate privileges to use the <code>listen()</code> function.</p>																
<b>RETURN VALUES</b>	Upon successful completions, <code>listen()</code> returns 0. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.																
<b>ERRORS</b>	<p>The <code>listen()</code> function will fail if:</p> <table><tr><td>EBADF</td><td>The <i>socket</i> argument is not a valid file descriptor.</td></tr><tr><td>EDESTADDRREQ</td><td>The socket is not bound to a local address, and the protocol does not support listening on an unbound socket.</td></tr><tr><td>EINVAL</td><td>The <i>socket</i> is already connected.</td></tr><tr><td>ENOTSOCK</td><td>The <i>socket</i> argument does not refer to a socket.</td></tr><tr><td>EOPNOTSUPP</td><td>The socket protocol does not support <code>listen()</code>.</td></tr></table> <p>The <code>listen()</code> function may fail if:</p> <table><tr><td>EACCES</td><td>The calling process does not have the appropriate privileges.</td></tr><tr><td>EINVAL</td><td>The <i>socket</i> has been shut down.</td></tr><tr><td>ENOBUFS</td><td>Insufficient resources are available in the system to complete the call.</td></tr></table>	EBADF	The <i>socket</i> argument is not a valid file descriptor.	EDESTADDRREQ	The socket is not bound to a local address, and the protocol does not support listening on an unbound socket.	EINVAL	The <i>socket</i> is already connected.	ENOTSOCK	The <i>socket</i> argument does not refer to a socket.	EOPNOTSUPP	The socket protocol does not support <code>listen()</code> .	EACCES	The calling process does not have the appropriate privileges.	EINVAL	The <i>socket</i> has been shut down.	ENOBUFS	Insufficient resources are available in the system to complete the call.
EBADF	The <i>socket</i> argument is not a valid file descriptor.																
EDESTADDRREQ	The socket is not bound to a local address, and the protocol does not support listening on an unbound socket.																
EINVAL	The <i>socket</i> is already connected.																
ENOTSOCK	The <i>socket</i> argument does not refer to a socket.																
EOPNOTSUPP	The socket protocol does not support <code>listen()</code> .																
EACCES	The calling process does not have the appropriate privileges.																
EINVAL	The <i>socket</i> has been shut down.																
ENOBUFS	Insufficient resources are available in the system to complete the call.																
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:																

listen(3XNET)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** accept(3XNET), connect(3XNET), socket(3XNET), attributes(5)

## netdir(3NSL)

<b>NAME</b>	netdir, netdir_getbyname, netdir_getbyaddr, netdir_free, netdir_options, taddr2uaddr, uaddr2taddr, netdir_perror, netdir_sperror, netdir_mergeaddr – generic transport name-to-address translation		
<b>SYNOPSIS</b>	<pre>#include &lt;netdir.h&gt;  int netdir_getbyname(const struct netconfig *config, const struct     nd_hostserv *service, struct nd_addrlist **addrs);  int netdir_getbyaddr(const struct netconfig *config, struct     nd_hostservlist **service, const struct netbuf *netaddr);  void netdir_free(void *ptr, const int struct_type);  int netdir_options(const struct netconfig *config, const int option,     const int fildes, char *point_to_args);  char *taddr2uaddr(const struct netconfig *config, const struct     netbuf *addr);  struct netbuf *uaddr2taddr(const struct netconfig *config, const     char *uaddr);  void netdir_perror(char *s);  char *netdir_sperror(void);</pre>		
<b>DESCRIPTION</b>	<p>The netdir routines provide a generic interface for name-to-address mapping that will work with all transport protocols. This interface provides a generic way for programs to convert transport specific addresses into common structures and back again. The netconfig structure, described on the netconfig(4) manual page, identifies the transport.</p> <p>The netdir_getbyname() routine maps the machine name and service name in the nd_hostserv structure to a collection of addresses of the type understood by the transport identified in the netconfig structure. This routine returns all addresses that are valid for that transport in the nd_addrlist structure. The nd_hostserv structure contains the following members:</p> <pre>char      /* host name */ *h_serv;  /* service name */</pre> <p>The nd_addrlist structure contains the following members:</p> <pre>int n_cnt;      /* number of addresses */ struct netbuf *n_addrs;</pre> <p>netdir_getbyname() accepts some special-case host names. The host names are defined in &lt;netdir.h&gt;. The currently defined host names are:</p> <table><tr><td>HOST_SELF</td><td>Represents the address to which local programs will bind their endpoints. HOST_SELF differs from the host name provided by gethostname(3C), which represents the address to which remote programs will bind their endpoints.</td></tr></table>	HOST_SELF	Represents the address to which local programs will bind their endpoints. HOST_SELF differs from the host name provided by gethostname(3C), which represents the address to which remote programs will bind their endpoints.
HOST_SELF	Represents the address to which local programs will bind their endpoints. HOST_SELF differs from the host name provided by gethostname(3C), which represents the address to which remote programs will bind their endpoints.		



HOST_ANY	Represents any host accessible by this transport provider. HOST_ANY allows applications to specify a required service without specifying a particular host name.
HOST_SELF_CONNECT	Represents the host address that can be used to connect to the local host.
HOST_BROADCAST	Represents the address for all hosts accessible by this transport provider. Network requests to this address are received by all machines.

All fields of the `nd_hostserv` structure must be initialized.

To find the address of a given host and service on all available transports, call the `netdir_getbyname()` routine with each `struct netconfig` structure returned by `getnetconfig(3NSL)`.

The `netdir_getbyaddr()` routine maps addresses to service names. This routine returns *service*, a list of host and service pairs that would yield this address. If more than one tuple of host and service name is returned, then the first tuple contains the preferred host and service names:

```
struct nd_hostservlist {
    int *h_cnt; /* number of hostservs found */
    struct hostserv *h_hostservs;
}
```

The `netdir_free()` structure is used to free the structures allocated by the name to address translation routines. *ptr* points to the structure that has to be freed. The `struct_type` identifies the structure:

```
struct netbuf          ND_ADDR
struct nd_addrlist    ND_ADDRLIST
struct hostserv       ND_HOSTSERV
struct nd_hostservlist ND_HOSTSERVLIST
```

Free the universal address returned by `taddr2uaddr()` by using `free()`.

The `netdir_options()` routine is used to do all transport-specific setups and option management. *fildev* is the associated file descriptor. *option*, *fildev*, and *pointer\_to\_args* are passed to the `netdir_options()` routine for the transport specified in *config*. Currently four values are defined for *option*:

```
ND_SET_BROADCAST
ND_SET_RESERVEDPORT
ND_CHECK_RESERVEDPORT
ND_MERGEADDR
```

## netdir(3NSL)

The `taddr2uaddr()` and `uaddr2taddr()` routines support translation between universal addresses and TLI type `netbufs`. The `taddr2uaddr()` routine takes a `struct netbuf` data structure and returns a pointer to a string that contains the universal address. It returns `NULL` if the conversion is not possible. This is not a fatal condition as some transports do not support a universal address form.

`uaddr2taddr()` is the reverse of `taddr2uaddr()`. It returns the `struct netbuf` data structure for the given universal address.

If a transport provider does not support an option, `netdir_options` returns `-1` and the error message can be printed through `netdir_perror()` or `netdir_serror()`.

The specific actions of each option follow.

### ND\_SET\_BROADCAST

Sets the transport provider up to allow broadcast, if the transport supports broadcast. *fildev* is a file descriptor into the transport, that is, the result of a `tc_open` of `/dev/udp`. *pointer\_to\_args* is not used. If this completes, broadcast operations can be performed on file descriptor *fildev*.

### ND\_SET\_RESERVEDPORT

Allows the application to bind to a reserved port, if that concept exists for the transport provider. *fildev* is an unbound file descriptor into the transport. If *pointer\_to\_args* is `NULL`, *fildev* is bound to a reserved port. If *pointer\_to\_args* is a pointer to a `netbuf` structure, an attempt is made to bind to any reserved port on the specified address.

### ND\_CHECK\_RESERVEDPORT

Used to verify that the address corresponds to a reserved port, if that concept exists for the transport provider. *fildev* is not used. *pointer\_to\_args* is a pointer to a `netbuf` structure that contains the address. This option returns `0` only if the address specified in *pointer\_to\_args* is reserved.

### ND\_MERGEADDR

Used to take a "local address" (like the `0.0.0.0` address that TCP uses) and return a "real address" that client machines can connect to. *fildev* is not used.

*pointer\_to\_args* is a pointer to a `struct nd_mergearg`, which has the following members:

```
char s_uaddr; /* server's universal address */
char c_uaddr; /* client's universal address */
char m_uaddr; /* the result */
```

If *s\_uaddr* is something like `0.0.0.0.1.12`, and, if the call is successful, *m\_uaddr* is set to something like `192.11.109.89.1.12`. For most transports, *m\_uaddr* is exactly what *s\_uaddr* is.

## RETURN VALUES

The `netdir_perror()` routine prints an error message on the standard output stating why one of the name-to-address mapping routines failed. The error message is preceded by the string given as an argument.

The `netdir_serror()` routine returns a string containing an error message stating why one of the name-to-address mapping routines failed.

`netdir_serror()` returns a pointer to a buffer which contains the error message string. This buffer is overwritten on each call. In multithreaded applications, this buffer is implemented as thread-specific data.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `gethostname(3C)`, `getnetconfig(3NSL)`, `getnetpath(3NSL)`, `netconfig(4)`, `attributes(5)`

## nis\_error(3NSL)

<b>NAME</b>	nis_error, nis_sperrno, nis_perror, nis_terror, nis_sperror, nis_sperror_r – display NIS+ error messages				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpcsvc/nis.h&gt;  char *nis_sperrno(nis_error status);  void nis_perror(nis_error status, char *label);  void nis_terror(nis_error status, char *label);  char *nis_sperror_r(nis_error status, char *label, char *buf, int length);  char *nis_sperror(nis_error status, char *label);</pre>				
<b>DESCRIPTION</b>	<p>These functions convert NIS+ status values into text strings.</p> <p>nis_sperrno() simply returns a pointer to a string constant which is the error string.</p> <p>nis_perror() prints the error message corresponding to <i>status</i> as “<i>label: error message</i>” on standard error.</p> <p>nis_terror() sends the error text to syslog(3C) at level LOG_ERR.</p> <p>The function nis_sperror_r(), returns a pointer to a string that can be used or copied using the strdup() function (See string(3C)). The caller must supply a string buffer, <i>buf</i>, large enough to hold the error string (a buffer size of 128 bytes is guaranteed to be sufficiently large). <i>status</i> and <i>label</i> are the same as for nis_perror(). The pointer returned by nis_sperror_r() is the same as <i>buf</i>, that is, the pointer returned by the function is a pointer to <i>buf</i>. <i>length</i> specifies the number of characters to copy from the error string to <i>buf</i>.</p> <p>The last function, nis_sperror(), is similar to nis_sperror_r() except that the string is returned as a pointer to a buffer that is reused on each call. nis_sperror_r() is the preferred interface, since it is suitable for single-threaded and multi-threaded programs.</p> <p>When compiling multithreaded applications, see Intro(3), Notes On Multithread Applications, for information about the use of the _REENTRANT flag.</p>				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
<b>SEE ALSO</b>	niserror(1), string(3C), syslog(3C), attributes(5)				

**NOTES** | NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <http://www.sun.com/directory/nisplus/transition.html>.

## nis\_groups(3NSL)

<b>NAME</b>	nis_groups, nis_ismember, nis_addmember, nis_removemember, nis_creategroup, nis_destroygroup, nis_verifygroup, nis_print_group_entry – NIS+ group manipulation functions
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpcsvc/nis.h&gt;  bool_t nis_ismember(nis_name principal, nis_name group); nis_error nis_addmember(nis_name member, nis_name group); nis_error nis_removemember(nis_name member, nis_name group); nis_error nis_creategroup(nis_name group, uint_t flags); nis_error nis_destroygroup(nis_name group); void nis_print_group_entry(nis_name group); nis_error nis_verifygroup(nis_name group);</pre>
<b>DESCRIPTION</b>	<p>These functions manipulate NIS+ groups. They are used by NIS+ clients and servers, and are the interfaces to the group authorization object.</p> <p>The names of NIS+ groups are syntactically similar to names of NIS+ objects but they occupy a separate namespace. A group named "a.b.c.d." is represented by a NIS+ group object named "a.groups_dir.b.c.d."; the functions described here all expect the name of the group, not the name of the corresponding group object.</p> <p>There are three types of group members:</p> <ul style="list-style-type: none"><li>■ An <i>explicit</i> member is just a NIS+ principal-name, for example "wickedwitch.west.oz."</li><li>■ An <i>implicit</i> ("domain") member, written "*.west.oz.", means that all principals in the given domain belong to this member. No other forms of wildcarding are allowed: "wickedwitch.*.oz." is invalid, as is "wickedwitch.west.*.". Note that principals in subdomains of the given domain are <i>not</i> included.</li><li>■ A <i>recursive</i> ("group") member, written "@cowards.oz.", refers to another group. All principals that belong to that group are considered to belong here.</li></ul> <p>Any member may be made <i>negative</i> by prefixing it with a minus sign ('-'). A group may thus contain explicit, implicit, recursive, negative explicit, negative implicit, and negative recursive members.</p> <p>A principal is considered to belong to a group if it belongs to at least one non-negative group member of the group and belongs to no negative group members.</p> <p>The <code>nis_ismember()</code> function returns TRUE if it can establish that <i>principal</i> belongs to <i>group</i>; otherwise it returns FALSE.</p> <p>The <code>nis_addmember()</code> and <code>nis_removemember()</code> functions add or remove a member. They do not check whether the member is valid. The user must have read and modify rights for the group in question.</p>

The `nis_creategroup()` and `nis_destroygroup()` functions create and destroy group objects. The user must have create or destroy rights, respectively, for the `groups_dir` directory in the appropriate domain. The parameter `flags` to `nis_creategroup()` is currently unused and should be set to zero.

The `nis_print_group_entry()` function lists a group's members on the standard output.

The `nis_verifygroup()` function returns `NIS_SUCCESS` if the given group exists, otherwise it returns an error code.

These functions only accept fully-qualified NIS+ names.

A group is represented by a NIS+ object with a variant part that is defined in the `group_obj` structure. See `nis_objects(3NSL)`. It contains the following fields:

```
uint_t    gr_flags;    /* Interpretation Flags
                       (currently unused) */
struct {
    uint_t    gr_members_len;
    nis_name  *gr_members_val;
} gr_members;    /* Array of members */
```

NIS+ servers and clients maintain a local cache of expanded groups to enhance their performance when checking for group membership. Should the membership of a group change, servers and clients with that group cached will not see the change until either the group cache has expired or it is explicitly flushed. A server's cache may be flushed programmatically by calling the `nis_servstate()` function with tag `TAG_GCACHE` and a value of 1.

There are currently no known methods for `nis_ismember()`, `nis_print_group_entry()`, and `nis_verifygroup()` to get their answers from only the master server.

## EXAMPLES

### EXAMPLE 1 Simple Memberships

Given a group `sadsouls.oz.` with members `tinman.oz.`, `lion.oz.`, and `scarecrow.oz.`, the function call

```
bool_var = nis_ismember("lion.oz.", "sadsouls.oz.");
```

will return 1 (TRUE) and the function call

```
bool_var = nis_ismember("toto.oz.", "sadsouls.oz.");
```

will return 0 (FALSE).

### EXAMPLE 2 Implicit Memberships

Given a group `baddies.oz.`, with members `wickedwitch.west.oz.` and `*.monkeys.west.oz.`, the function call `bool_var = nis_ismember("hogan.monkeys.west.oz.", "baddies.oz.");` will return 1 (TRUE) because any principal from the `monkeys.west.oz.` domain belongs to the implicit group `*.monkeys.west.oz.`, but the function call

## nis\_groups(3NSL)

### **EXAMPLE 2** Implicit Memberships *(Continued)*

```
bool_var = nis_ismember("hogan.big.monkeys.west.oz.", "baddies.oz.");  
will return 0 (FALSE).
```

### **EXAMPLE 3** Recursive Memberships

Given a group `goodandbad.oz.`, with members `toto.kansas`, `@sadsouls.oz.`, and `@baddies.oz.`, and the groups `sadsouls.oz.` and `baddies.oz.` defined above, the function call

```
bool_var = nis_ismember("wickedwitch.west.oz.", "goodandbad.oz.");
```

will return 1 (TRUE), because `wickedwitch.west.oz.` is a member of the `baddies.oz.` group which is recursively included in the `goodandbad.oz.` group.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `nisrpadm(1)`, `nis_objects(3NSL)`, `attributes(5)`

**NOTES** NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <http://www.sun.com/directory/nisplus/transition.html>.



<b>NAME</b>	nisl_local_names, nisl_local_directory, nisl_local_host, nisl_local_group, nisl_local_principal – NIS+ local names				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpcsvc/nis.h&gt;  nis_name nisl_local_directory(void); nis_name nisl_local_host(void); nis_name nisl_local_group(void); nis_name nisl_local_principal(void);</pre>				
<b>DESCRIPTION</b>	<p>These functions return several default NIS+ names associated with the current process.</p> <p><code>nisl_local_directory()</code> returns the name of the NIS+ domain for this machine. This is currently the same as the Secure RPC domain returned by the <code>sysinfo(2)</code> system call.</p> <p><code>nisl_local_host()</code> returns the NIS+ name of the current machine. This is the fully qualified name for the host and is either the value returned by the <code>gethostname(3C)</code> function or, if the host name is only partially qualified, the concatenation of that value and the name of the NIS+ directory. Note that if a machine's name and address cannot be found in the local NIS+ directory, its hostname must be fully qualified.</p> <p><code>nisl_local_group()</code> returns the name of the current NIS+ group name. This is currently set by setting the environment variable <code>NIS_GROUP</code> to the groupname.</p> <p><code>nisl_local_principal()</code> returns the NIS+ principal name for the user associated with the effective UID of the calling process. This function maps the effective uid into a principal name by looking for a LOCAL type credential in the table named <code>cred.org_dir</code> in the default domain. See <code>nisladdcred(1M)</code>.</p> <p>The result returned by these routines is a pointer to a data structure with the NIS+ library, and should be considered a "read-only" result and should not be modified.</p>				
<b>ENVIRONMENT VARIABLES</b>	<p><code>NIS_GROUP</code> This variable contains the name of the local NIS+ group. If the name is not fully qualified, the value returned by <code>nisl_local_directory()</code> will be concatenated to it.</p>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>nisldefaults(1)</code> , <code>nisladdcred(1M)</code> , <code>sysinfo(2)</code> , <code>gethostname(3C)</code> , <code>nisl_names(3NSL)</code> , <code>nisl_objects(3NSL)</code> , <code>attributes(5)</code>				

nis\_local\_names(3NSL)

**NOTES** | NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <http://www.sun.com/directory/nisplus/transition.html>.

<b>NAME</b>	nis_names, nis_lookup, nis_add, nis_remove, nis_modify, nis_freeresult – NIS+ namespace functions		
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpcsvc/nis.h&gt;  nis_result *nis_lookup(nis_name name, uint_t flags); nis_result *nis_add(nis_name name, nis_object *obj); nis_result *nis_remove(nis_name name, nis_object *obj); nis_result *nis_modify(nis_name name, nis_object *obj); void nis_freeresult(nis_result *result);</pre>		
<b>DESCRIPTION</b>	<p>The NIS+ namespace functions are used to locate and manipulate all NIS+ objects except the NIS+ entry objects. See <code>nis_objects(3NSL)</code>. To look up the NIS+ entry objects within a NIS+ table, refer to <code>nis_subr(3NSL)</code>.</p> <p><code>nis_lookup()</code> resolves a NIS+ name and returns a copy of that object from a NIS+ server. <code>nis_add()</code> and <code>nis_remove()</code> add and remove objects to the NIS+ namespace, respectively. <code>nis_modify()</code> can change specific attributes of an object that already exists in the namespace.</p> <p>These functions should be used only with names that refer to an NIS+ Directory, NIS+ Table, NIS+ Group, or NIS+ Private object. If a name refers to an NIS+ entry object, the functions listed in <code>nis_subr(3NSL)</code> should be used.</p> <p><code>nis_freeresult()</code> frees all memory associated with a <code>nis_result</code> structure. This function must be called to free the memory associated with a NIS+ result.</p> <p><code>nis_lookup()</code>, <code>nis_add()</code>, <code>nis_remove()</code>, and <code>nis_modify()</code> all return a pointer to a <code>nis_result()</code> structure which must be freed by calling <code>nis_freeresult()</code> when you have finished using it. If one or more of the objects returned in the structure need to be retained, they can be copied with <code>nis_clone_object(3NSL)</code>. See <code>nis_subr(3NSL)</code>.</p> <p><code>nis_lookup()</code> takes two parameters, the name of the object to be resolved in <i>name</i>, and a flags parameter, <i>flags</i>, which is defined below. The object name is expected to correspond to the syntax of a non-indexed NIS+ name. See <code>nis_tables(3NSL)</code>. The <code>nis_lookup()</code> function is the only function from this group that can use a non-fully qualified name. If the parameter <i>name</i> is not a fully qualified name, then the flag <code>EXPAND_NAME</code> must be specified in the call. If this flag is not specified, the function will fail with the error <code>NIS_BADNAME</code>.</p> <p>The <i>flags</i> parameter is constructed by logically ORing zero or more flags from the following list.</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>FOLLOW_LINKS</code></td> <td>When specified, the client library will “follow” links by issuing another NIS+ lookup call for the object named by the link. If the linked object is itself a link, then this</td> </tr> </table>	<code>FOLLOW_LINKS</code>	When specified, the client library will “follow” links by issuing another NIS+ lookup call for the object named by the link. If the linked object is itself a link, then this
<code>FOLLOW_LINKS</code>	When specified, the client library will “follow” links by issuing another NIS+ lookup call for the object named by the link. If the linked object is itself a link, then this		

## nis\_names(3NSL)

	process will iterate until the either a object is found that is not a LINK type object, or the library has followed 16 links.
HARD_LOOKUP	When specified, the client library will retry the lookup until it is answered by a server. Using this flag will cause the library to block until at least one NIS+ server is available. If the network connectivity is impaired, this can be a relatively long time.
NO_CACHE	When specified, the client library will bypass any object caches and will get the object from either the master NIS+ server or one of its replicas.
MASTER_ONLY	When specified, the client library will bypass any object caches and any domain replicas and fetch the object from the NIS+ master server for the object's domain. This insures that the object returned is up to date at the cost of a possible performance degradation and failure if the master server is unavailable or physically distant.
EXPAND_NAME	When specified, the client library will attempt to expand a partially qualified name by calling the function <code>nis_getnames()</code> , which uses the environment variable <code>NIS_PATH</code> . See <code>nis_subr(3NSL)</code> .

The status value may be translated to ASCII text using the function `nis_sperrno()`. See `nis_error(3NSL)`.

On return, the *objects* array in the result will contain one and possibly several objects that were resolved by the request. If the `FOLLOW_LINKS` flag was present, on success the function could return several entry objects if the link in question pointed within a table. If an error occurred when following a link, the objects array will contain a copy of the link object itself.

The function `nis_add()` will take the object *obj* and add it to the NIS+ namespace with the name *name*. This operation will fail if the client making the request does not have the *create* access right for the domain in which this object will be added. The parameter *name* must contain a fully qualified NIS+ name. The object members *zo\_name* and *zo\_domain* will be constructed from this name. This operation will fail if the object already exists. This feature prevents the accidental addition of objects over another object that has been added by another process.

The function `nis_remove()` will remove the object with name *name* from the NIS+ namespace. The client making this request must have the *destroy* access right for the domain in which this object resides. If the named object is a link, the link is removed and *not* the object that it points to. If the parameter *obj* is not `NULL`, it is assumed to point to a copy of the object being removed. In this case, if the object on the server

does not have the same object identifier as the object being passed, the operation will fail with the NIS\_NOTSAMEOBJ error. This feature allows the client to insure that it is removing the desired object. The parameter *name* must contain a fully qualified NIS+ name.

The function `nis_modify()` will modify the object named by *name* to the field values in the object pointed to by *obj*. This object should contain a copy of the object from the name space that is being modified. This operation will fail with the error NIS\_NOTSAMEOBJ if the object identifier of the passed object does not match that of the object being modified in the namespace.

Normally the contents of the member *zo\_name* in the *nis\_object* structure would be constructed from the name passed in the *name* parameter. However, if it is non-null the client library will use the name in the *zo\_name* member to perform a rename operation on the object. This name *must not* contain any unquoted '.'(dot) characters. If these conditions are not met the operation will fail and return the NIS\_BADNAME error code.

You cannot modify the name of an object if that modification would cause the object to reside in a different domain.

You cannot modify the schema of a table object.

**Results** These functions return a pointer to a structure of type `nis_result`:

```
struct nis_result {
    nis_error status;
    struct {
        uint_t    objects_len;
        nis_object *objects_val;
    } objects;
    netobj    cookie;
    uint32_t  zticks;
    uint32_t  dticks;
    uint32_t  aticks;
    uint32_t  cticks;
};
```

The *status* member contains the error status of the the operation. A text message that describes the error can be obtained by calling the function `nis_sperrno()`. See `nis_error(3NSL)`.

The *objects* structure contains two members. *objects\_val* is an array of *nis\_object* structures; *objects\_len* is the number of cells in the array. These objects will be freed by the call to `nis_freeresult()`. If you need to keep a copy of one or more objects, they can be copied with the function `nis_clone_object()` and freed with the function `nis_destroy_object()`. See `nis_server(3NSL)`. Refer to `nis_objects(3NSL)` for a description of the *nis\_object* structure.

## nis\_names(3NSL)

The various ticks contain details of where the time was taken during a request. They can be used to tune one's data organization for faster access and to compare different database implementations.

<i>zticks</i>	The time spent in the NIS+ service itself. This count starts when the server receives the request and stops when it sends the reply.
<i>dticks</i>	The time spent in the database backend. This time is measured from the time a database call starts, until the result is returned. If the request results in multiple calls to the database, this is the sum of all the time spent in those calls.
<i>aticks</i>	The time spent in any "accelerators" or caches. This includes the time required to locate the server needed to resolve the request.
<i>cticks</i>	The total time spent in the request. This clock starts when you enter the client library and stops when a result is returned. By subtracting the sum of the other ticks values from this value, you can obtain the local overhead of generating a NIS+ request.

Subtracting the value in *dticks* from the value in *zticks* will yield the time spent in the service code itself. Subtracting the sum of the values in *zticks* and *aticks* from the value in *cticks* will yield the time spent in the client library itself. Note: all of the tick times are measured in microseconds.

## RETURN VALUES

The client library can return a variety of error returns and diagnostics. The more salient ones are documented below.

NIS_SUCCESS	The request was successful.
NIS_S_SUCCESS	The request was successful, however the object returned came from an object cache and not directly from the server. If you do not wish to see objects from object caches you must specify the flag NO_CACHE when you call the lookup function.
NIS_NOTFOUND	The named object does not exist in the namespace.
NIS_CACHEEXPIRED	The object returned came from an object cache taht has <i>expired</i> . The time to live value has gone to zero and the object may have changed. If the flag NO_CACHE was passed to the lookup function then the lookup function will retry the operation to get an unexpired copy of the object.
NIS_NAMEUNREACHABLE	A server for the directory of the named object could not be reached. This can occur

	when there is a network partition or all servers have crashed. See the <code>HARD_LOOKUP</code> flag.
<code>NIS_UNKNOWNOBJ</code>	The object returned is of an unknown type.
<code>NIS_TRYAGAIN</code>	The server connected to was too busy to handle your request. For the <i>add</i> , <i>remove</i> , and <i>modify</i> operations this is returned when either the master server for a directory is unavailable, or it is in the process of checkpointing its database. It can also be returned when the server is updating its internal state. In the case of <code>nis_list()</code> , <code>NIS_TRYAGAIN</code> is returned if the client specifies a callback and the server does not have enough resources to handle the callback.
<code>NIS_SYSTEMERROR</code>	A generic system error occurred while attempting the request. Most commonly the server has crashed or the database has become corrupted. Check the syslog record for error messages from the server.
<code>NIS_NOT_ME</code>	A request was made to a server that does not serve the name in question. Normally this will not occur, however if you are not using the built in location mechanism for servers you may see this if your mechanism is broken.
<code>NIS_NOMEMORY</code>	Generally a fatal result. It means that the service ran out of heap space.
<code>NIS_NAMEEXISTS</code>	An attempt was made to add a name that already exists. To add the name, first remove the existing name and then add the new object or modify the existing named object.
<code>NIS_NOTMASTER</code>	An attempt was made to update the database on a replica server.
<code>NIS_INVALIDOBJ</code>	The object pointed to by <i>obj</i> is not a valid NIS+ object.
<code>NIS_BADNAME</code>	The name passed to the function is not a legal NIS+ name.

## nis\_names(3NSL)

NIS_LINKNAMEERROR	The name passed resolved to a LINK type object and the contents of the link pointed to an invalid name.
NIS_NOTSAMEOBJ	An attempt to remove an object from the namespace was aborted because the object that would have been removed was not the same object that was passed in the request.
NIS_NOSUCHNAME	This hard error indicates that the named directory of the table object does not exist. This occurs when the server that should be the parent of the server that serves the table, does not know about the directory in which the table resides.
NIS_NOSUCHTABLE	The named table does not exist.
NIS_MODFAIL	The attempted modification failed.
NIS_FOREIGNNS	The name could not be completely resolved. When the name passed to the function would resolve in a namespace that is outside the NIS+ name tree, this error is returned with a NIS+ object of type DIRECTORY, which contains the type of namespace and contact information for a server within that namespace.
NIS_RPCERROR	This fatal error indicates the RPC subsystem failed in some way. Generally there will be a syslog(3C) message indicating why the RPC request failed.

### ENVIRONMENT VARIABLES

NIS\_PATH If the flag EXPAND\_NAME is set, this variable is the search path used by nis\_lookup().

### ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

### SEE ALSO

nis\_error(3NSL), nis\_objects(3NSL), nis\_server(3NSL), nis\_subr(3NSL), nis\_tables(3NSL), attributes(5)

### NOTES

NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <http://www.sun.com/directory/nisplus/transition.html>.



<b>NAME</b>	nisl_objects – NIS+ object formats
<b>SYNOPSIS</b>	cc [ <i>flag ...</i> ] <i>file ...</i> -lnsl [ <i>library ...</i> ] /usr/include/rpcsvc/nisl_objects.x
<b>Common Attributes</b>	<p>The NIS+ service uses a variant record structure to hold the contents of the objects that are used by the NIS+ service. These objects all share a common structure that defines a set of attributes that all objects possess. The <code>nisl_object</code> structure contains the following members:</p> <pre>typedef  char    *nisl_name;         struct  nisl_object {             nisl_oid      zo_oid;             nisl_name     zo_name;             nisl_name     zo_owner;             nisl_name     zo_group;             nisl_name     zo_domain;             uint_t        zo_access;             uint32_t      zo_ttl;             objdata      zo_data;         };</pre> <p>In this structure, the first member <code>zo_oid</code>, is a 64 bit number that uniquely identifies this instance of the object on this server. This member is filled in by the server when the object is created and changed by the server when the object is modified. When used in conjunction with the object's name and domain it uniquely identifies the object in the entire NIS+ namespace.</p> <p>The second member, <code>zo_name</code>, contains the leaf name of the object. This name is never terminated with a '.' (dot). When an object is created or added to the namespace, the client library will automatically fill in this field and the domain name from the name that was passed to the function.</p> <p><code>zo_domain</code> contains the name of the NIS+ domain to which this object belongs. This information is useful when tracking the parentage of an object from a cache. When used in conjunction with the members <code>zo_name</code> and <code>zo_oid</code>, it uniquely identifies an object. This makes it possible to always reconstruct the name of an object by using the code fragment</p> <pre>printf(buf, "%s.%s", obj-&gt;zo_name, obj-&gt;zo_domain);</pre> <p>The <code>zo_owner</code> and <code>zo_group</code> members contain the NIS+ names of the object's principal owner and group owner, respectively. Both names must be NIS+ fully qualified names. However, neither name can be used directly to identify the object they represent. This stems from the condition that NIS+ uses itself to store information that it exports.</p> <p>The <code>zo_owner</code> member contains a fully qualified NIS+ name of the form <i>principal.domain</i>. This name is called a NIS+ principal name and is used to identify authentication information in a credential table. When the server constructs a search query of the form</p> <pre>[cname=principal], cred.org_dir.domain.</pre>

## nis\_objects(3NSL)

The query will return to the server credential information about *principal* for all flavors of RPC authentication that are in use by that principal. When an RPC request is made to the server, the authentication flavor is extracted from the request and is used to find out the NIS+ principal name of the client. For example, if the client is using the AUTH\_DES authentication flavor, it will include in the authentication credentials the network name or *netname* of the user making the request. This netname will be of the form

*unix.UID@domain*

The NIS+ server will then construct a query on the credential database of the form

`[auth_name=netname, auth_type=AUTH_DES], cred.org_dir.domain.`

This query will return an entry which contains a principal name in the first column. This NIS+ principal name is used to control access to NIS+ objects.

The group owner for the object is treated differently. The group owner member is optional (it should be the null string if not present) but must be fully qualified if present. A group name takes the form

*group.domain.*

which the server then maps into a name of the form

*group.groups\_dir.*domain*.*

The purpose of this mapping is to prevent NIS+ group names from conflicting with user specified domain or table names. For example, if a domain was called *engineering.foo.com.*, then without the mapping a NIS+ group of the same name to represent members of engineering would not be possible. The contents of groups are lists of NIS+ principal names which are used exactly like the `zo_owner` name in the object. See `nis_groups(3NSL)` for more details.

The `zo_access` member contains the bitmask of access rights assigned to this object. There are four access rights defined, and four are reserved for future use and must be zero. This group of 8 access rights can be granted to four categories of client. These categories are the object's owner, the object's group owner, all authenticated clients (world), and all unauthenticated clients (nobody). Note that access granted to "nobody" is really access granted to everyone, authenticated and unauthenticated clients.

The `zo_ttl` member contains the number of seconds that the object can "live" in a cache before it is expired. This value is called the time to live for this object. This number is particularly important on group and directory (domain) objects. When an object is cached, the current time is added to the value in `zo_ttl`. Then each time the cached object is used, the time in `zo_ttl` is compared with the current time. If the current time is later than the time in `zo_ttl` the object is said to have expired and the cached copy should not be used.

Setting the TTL is somewhat of an art. You can think of it as the “half life” of the object, or half the amount of time you believe will pass before the object changes. The benefit of setting the ttl to a large number is that the object will stay in a cache for long periods of time. The problem with setting it to a large value is that when the object changes it will take a long time for the caches to flush out old copies of that object. The problems and benefits are reversed for setting the time to a small value. Generally setting the value to 43200 (12 hrs) is reasonable for things that change day to day, and 3024000 is good for things that change week to week. Setting the value to 0 will prevent the object from ever being cached since it would expire immediately.

The `zo_data` member is a discriminated union with the following members:

```
zotypes zo_type;
union {
    struct directory_obj    di_data;
    struct group_obj       gr_data;
    struct table_obj       ta_data;
    struct entry_obj       en_data;
    struct link_obj        li_data;
    struct {
        uint_t    po_data_len;
        char      *po_data_val;
    } po_data;
} objdata_u;
```

The union is discriminated based on the type value contained in `zo_type`. There six types of objects currently defined in the NIS+ service. These types are the directory, link, group, table, entry, and private types.

```
enum zotypes {
    BOGUS_OBJ    = 0,
    NO_OBJ       = 1,
    DIRECTORY_OBJ = 2,
    GROUP_OBJ    = 3,
    TABLE_OBJ   = 4,
    ENTRY_OBJ    = 5,
    LINK_OBJ     = 6,
    PRIVATE_OBJ  = 7
};
typedef enum zotypes zotypes;
```

All object types define a structure that contains data specific to that type of object. The simplest are private objects which are defined to contain a variable length array of octets. Only the owner of the object is expected to understand the contents of a private object. The following section describe the other five object types in more significant detail.

## Directory Objects

The first type of object is the *directory* object. This object’s variant part is defined as follows:

```
enum nstype {
    UNKNOWN    = 0,
    NIS        = 1,
    SUNYP      = 2,
```

## nis\_objects(3NSL)

```
    DNS      = 4,
    X500     = 5,
    DNANS    = 6,
    XCHS     = 7,
}
typedef enum nstype nstype;
struct oar_mask {
    uint_t    oa_rights;
    zotypes   oa_otype;
}
typedef struct oar_mask oar_mask;
struct endpoint {
    char      *uaddr;
    char      *family;
    char      *proto;
}
typedef struct endpoint endpoint;
struct nis_server {
    nis_name   name;
    struct {
        uint_t    ep_len;
        endpoint   *ep_val;
    } ep;
    uint_t    key_type;
    netobj    pkey;
}
typedef struct nis_server nis_server;
struct directory_obj {
    nis_name   do_name;
    nstype     do_type;
    struct {
        uint_t    do_servers_len;
        nis_server *do_servers_val;
    } do_servers;
    uint32_t    do_ttl;
    struct {
        uint_t    do_armask_len;
        oar_mask  *do_armask_val;
    } do_armask;
}
typedef struct directory_obj directory_obj;
```

The main structure contains five primary members: `do_name`, `do_type`, `do_servers`, `do_ttl`, and `do_armask`. The information in the `do_servers` structure is sufficient for the client library to create a network connection with the named server for the directory.

The `do_name` member contains the name of the directory or domain represented in a format that is understandable by the type of nameservice serving that domain. In the case of NIS+ domains, this is the same as the name that can be composed using the `zo_name` and `zo_domain` members. For other name services, this name will be a name that they understand. For example, if this were a directory object describing an X.500 namespace that is "under" the NIS+ directory *eng.sun.com.*, this name might contain `"/C=US, /O=Sun Microsystems, /OU=Engineering/"`. The type of nameservice that is being described is determined by the value of the member `do_type`.

The `do_servers` structure contains two members. `do_servers_val` is an array of `nis_server` structures; `do_servers_len` is the number of cells in the array. The `nis_server` structure is designed to contain enough information such that machines on the network providing name services can be contacted without having to use a name service. In the case of NIS+ servers, this information is the name of the machine in `name`, its public key for authentication in `pkey`, and a variable length array of endpoints, each of which describes the network endpoint for the `rpcbind` daemon on the named machine. The client library uses the addresses to contact the server using a transport that both the client and server can communicate on and then queries the `rpcbind` daemon to get the actual transport address that the server is using.

Note that the first server in the `do_servers` list is always the master server for the directory.

The `key_type` field describes the type of key stored in the `pkey` netobj (see `/usr/include/rpc/xdr.h` for a definition of the network object structure). Currently supported types are `NIS_PK_NONE` for no public key, `NIS_PK_DH` for a Diffie-Hellman type public key, and `NIS_PK_DHEXT` for an extended Diffie-Hellman public key.

The `do_ttl` member contains a copy of the `zo_ttl` member from the common attributes. This is duplicated because the cache manager only caches the variant part of the directory object.

The `do_armask` structure contains two members. `do_armask_val` is an array of `oar_mask` structures; `do_armask_len` is the number of cells in the array. The `oar_mask` structure contains two members: `oa_rights` specifies the access rights allowed for objects of type `oa_otype`. These access rights are used for objects of the given type in the directory when they are present in this array.

The granting of access rights for objects contained within a directory is actually two-tiered. If the directory object itself grants a given access right (using the `zo_access` member in the `nis_object` structure representing the directory), then all objects within the directory are allowed that access. Otherwise, the `do_armask` structure is examined to see if the access is allowed specifically for that type of structure. This allows the administrator of a namespace to set separate policies for different object types, for example, one policy for the creation of tables and another policy for the creation of other directories. See `nis+(1)` for more details.

## Link Objects

Link objects provide a means of providing *aliases* or symbolic links within the namespace. Their variant part is defined as follows.

```
struct link_obj {
    zotypes    li_rtype;
    struct {
        uint_t    li_attrs_len;
        nis_attr  *li_attrs_val;
    } li_attrs;
    nis_name li_name;
}
```

## nis\_objects(3NSL)

The `li_rtype` member contains the object type of the object pointed to by the link. This is only a hint, since the object which the link points to may have changed or been removed. The fully qualified name of the object (table or otherwise) is specified in the member `li_name`.

NIS+ links can point to either other objects within the NIS+ namespace, or to entries within a NIS+ table. If the object pointed to by the link is a table and the member `li_attrs` has a nonzero number of attributes (index name/value pairs) specified, the table is searched when this link is followed. All entries which match the specified search pattern are returned. Note, that unless the flag `FOLLOW_LINKS` is specified, the `nis_lookup(3NSL)` function will always return non-entry objects.

### Group Objects

Group objects contain a membership list of NIS+ principals. The group objects' variant part is defined as follows.

```
struct group_obj {
    uint_t    gr_flags;
    struct {
        uint_t    gr_members_len;
        nis_name *gr_members_val;
    } gr_members;
}
```

The `gr_flags` member contains flags that are currently unused. The `gr_members` structure contains the list of principals. For a complete description of how group objects are manipulated see `nis_groups(3NSL)`.

### Table Objects

The NIS+ table object is analogous to a YP map. The differences stem from the access controls, and the variable schemas that NIS+ allows. The table objects data structure is defined as follows:

```
#define TA_BINARY    1
#define TA_CRYPT    2
#define TA_XDR      4
#define TA_SEARCHABLE 8
#define TA_CASE     16
#define TA_MODIFIED 32
struct table_col {
    char    *tc_name;
    uint_t  tc_flags;
    uint_t  tc_rights;
}
typedef struct table_col table_col;
struct table_obj {
    char    *ta_type;
    uint_t  ta_maxcol;
    uchar_t ta_sep;
    struct {
        uint_t  ta_cols_len;
        table_col *ta_cols_val;
    } ta_cols;
    char    *ta_path;
}
```

The `ta_type` member contains a string that identifies the type of entries in this table. NIS+ does not enforce any policies as to the contents of this string. However, when entries are added to the table, the NIS+ service will check to see that they have the same "type" as the table as specified by this member.

The structure `ta_cols` contains two members. `ta_cols_val` is an array of `table_col` structures. The length of the array depends on the number of columns in the table; it is defined when the table is created and is stored in `ta_cols_len`. `ta_maxcol` also contains the number of columns in the table and always has the same value as `ta_cols_len`. Once the table is created, this length field cannot be changed.

The `ta_sep` character is used by client applications that wish to print out an entry from the table. Typically this is either space (" ") or colon (":").

The `ta_path` string defines a concatenation path for tables. This string contains an ordered list of fully qualified table names, separated by colons, that are to be searched if a search on this table fails to match any entries. This path is only used with the flag `FOLLOW_PATH` with a `nis_list()` call. See `nis_tables(3NSL)` for information on these flags.

In addition to checking the type, the service will check that the number of columns in an entry is the same as those in the table before allowing that entry to be added.

Each column has associated with it a name in `tc_name`, a set of flags in `tc_flags`, and a set of access rights in `tc_rights`. The name should be indicative of the contents of that column.

The `TA_BINARY` flag indicates that data in the column is binary (rather than text). Columns that are searchable cannot contain binary data. The `TA_CRYPT` flag specifies that the information in this column should be encrypted prior to sending it over the network. This flag has no effect in the export version of NIS+. The `TA_XDR` flag is used to tell the client application that the data in this column is encoded using the XDR protocol. The `TA_BINARY` flag must be specified with the XDR flag. Further, by convention, the name of a column that has the `TA_XDR` flag set is the name of the XDR function that will decode the data in that column.

The `TA_SEARCHABLE` flag specifies that values in this column can be searched. Searchable columns must contain textual data and must have a name associated with them. The flag `TA_CASE` specifies that searches involving this column ignore the case of the value in the column. At least one of the columns in the table should be searchable. Also, the combination of all searchable column values should uniquely select an entry within the table. The `TA_MODIFIED` flag is set only when the table column is modified. When `TA_MODIFIED` is set, and the object is modified again, the modified access rights for the table column must be copied, not the default access rights.

## Entry Objects

Entry objects are stored in tables. The structure used to define the entry data is as follows.

## nis\_objects(3NSL)

```
#define EN_BINARY    1
#define EN_CRYPT    2
#define EN_XDR      4
#define EN_MODIFIED  8
struct entry_col {
    uint_t    ec_flags;
    struct {
        uint_t    ec_value_len;
        char      *ec_value_val;
    } ec_value;
}
typedef struct entry_col entry_col;
struct entry_obj {
    char      *en_type;
    struct {
        uint_t    en_cols_len;
        entry_col *en_cols_val;
    } en_cols;
}
```

The `en_type` member contains a string that specifies the type of data this entry represents. The NIS+ server will compare this string to the type string specified in the table object and disallow any updates or modifications if they differ.

The `en_cols` structure contains two members: `en_cols_len` and `en_cols_val`. `en_cols_val` is an array of `entry_col` structures. `en_cols_len` contains a count of the number of cells in the `en_cols_val` array and reflects the number of columns in the table -- it always contains the same value as the `table_obj.ta_cols.ta_cols_len` member from the table which contains the entry.

The `entry_col` structure contains information about the entry's per-column values. `ec_value` contains information about a particular value. It has two members: `ec_value_val`, which is the value itself, and `ec_value_len`, which is the length (in bytes) of the value. `entry_col` also contains the member `ec_flags`, which contains a set of flags for the entry.

The flags in `ec_flags` are primarily used when adding or modifying entries in a table. All columns that have the flag `EN_CRYPT` set will be encrypted prior to sending them over the network. Columns with `EN_BINARY` set are presumed to contain binary data. The server will ensure that the column in the table object specifies binary data prior to allowing the entry to be added. When modifying entries in a table, only those columns that have changed need be sent to the server. Those columns should each have the `EN_MODIFIED` flag set to indicate this to the server.

**SEE ALSO** `nis+(1)`, `nis_groups(3NSL)`, `nis_names(3NSL)`, `nis_server(3NSL)`, `nis_subr(3NSL)`, `nis_tables(3NSL)`

**NOTES** NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <http://www.sun.com/directory/nisplus/transition.html>.



<b>NAME</b>	nis_ping, nis_checkpoint – NIS+ log administration functions				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpcsvc/nis.h&gt;  void nis_ping(nis_name dirname, uint32_t utime, nis_object *diobj); nis_result *nis_checkpoint(nis_name dirname);</pre>				
<b>DESCRIPTION</b>	<p>nis_ping() is called by the master server for a directory when a change has occurred within that directory. The parameter <i>dirname</i> identifies the directory with the change. If the parameter <i>diobj</i> is NULL, this function looks up the directory object for <i>dirname</i> and uses the list of replicas it contains. The parameter <i>utime</i> contains the timestamp of the last change made to the directory. This timestamp is used by the replicas when retrieving updates made to the directory.</p> <p>The effect of calling nis_ping() is to schedule an update on the replica. A short time after a ping is received, typically about two minutes, the replica compares the last update time for its databases to the timestamp sent by the ping. If the ping timestamp is later, the replica establishes a connection with the master server and request all changes from the log that occurred after the last update that it had recorded in its local log.</p> <p>nis_checkpoint() is used to force the service to checkpoint information that has been entered in the log but has not been checkpointed to disk. When called, this function checkpoints the database for each table in the directory, the database containing the directory and the transaction log. Care should be used in calling this function since directories that have seen a lot of changes may take several minutes to checkpoint. During the checkpointing process, the service will be unavailable for updates for all directories that are served by this machine as master.</p> <p>nis_checkpoint() returns a pointer to a <i>nis_result</i> structure. See nis_tables(3NSL). This structure should be freed with nis_freeresult(). See nis_names(3NSL). The only items of interest in the returned result are the status value and the statistics.</p>				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	nislog(1M), nis_names(3NSL), nis_tables(3NSL), nisfiles(4), attributes(5)				
<b>NOTES</b>	NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <a href="http://www.sun.com/directory/nisplus/transition.html">http://www.sun.com/directory/nisplus/transition.html</a> .				

nis\_server(3NSL)

<b>NAME</b>	nis_server, nis_mkdir, nis_rmdir, nis_servstate, nis_stats, nis_getservlist, nis_freeservlist, nis_freetags – miscellaneous NIS+ functions
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpcsvc/nis.h&gt;  nis_error nis_mkdir(nis_name dirname, nis_server *machine); nis_error nis_rmdir(nis_name dirname, nis_server *machine); nis_error nis_servstate(nis_server *machine, nis_tag *tags, int     numtags, nis_tag **result); nis_error nis_stats(nis_server *machine, nis_tag *tags, int numtags,     nis_tag **result); void nis_freetags(nis_tag *tags, int numtags); nis_server **nis_getservlist(nis_name dirname); void nis_freeservlist(nis_server **machines);</pre>
<b>DESCRIPTION</b>	<p>These functions provide a variety of services for NIS+ applications.</p> <p>nis_mkdir() is used to create the necessary databases to support NIS+ service for a directory, <i>dirname</i>, on a server, <i>machine</i>. If this operation is successful, it means that the directory object describing <i>dirname</i> has been updated to reflect that server <i>machine</i> is serving the named directory. For a description of the nis_server structure, refer to nis_objects(3NSL).</p> <p>Per-server and per-directory access restrictions may apply to nis_mkdir(). See nisopaccess(1).</p> <p>nis_rmdir() is used to delete the directory, <i>dirname</i>, from the specified server machine. The <i>machine</i> parameter cannot be NULL. Note that nis_rmdir() does not remove the directory <i>dirname</i> from the namespace or remove a server from the server list in the directory object. To remove a directory from the namespace you must call nis_remove() to remove the directory <i>dirname</i> from the namespace and call nis_rmdir() for each server in the server list to remove the directory from the server. To remove a replica from the server list, you need to first call nis_modify() to remove the server from the directory object and then call nis_rmdir() to remove the replica.</p> <p>Per-server and per-directory access restrictions may apply to nis_rmdir(). See nisopaccess(1).</p> <p>For a description of the nis_server structure, refer to nis_objects(3NSL).</p> <p>nis_servstate() is used to set and read the various state variables of the NIS+ servers. In particular the internal debugging state of the servers may be set and queried.</p>

The `nis_stats()` function is used to retrieve statistics about how the server is operating. Tracking these statistics can help administrators determine when they need to add additional replicas or to break up a domain into two or more subdomains. For more information on reading statistics, see `nisstat(1M)`.

`nis_servstate()` and `nis_stats()` use the tag list. This tag list is a variable length array of `nis_tag` structures whose length is passed to the function in the `numtags` parameter. The set of legal tags are defined in the file `<rpcsvc/nis_tags.h>` which is included in `<rpcsvc/nis.h>`. Because these tags can and do vary between implementations of the NIS+ service, it is best to consult this file for the supported list. Passing unrecognized tags to a server will result in their `tag_value` member being set to the string "unknown." Both of these functions return their results in malloced tag structure, `*result`. If there is an error, `*result` is set to `NULL`. The `tag_value` pointers points to allocated string memory which contains the results. Use `nis_frehtags()` to free the tag structure.

Per-server and per-directory access restrictions may apply to the `NIS_SERVSTATE` or `NIS_STATUS` (`nis_stats()`) operations and their sub-operations (`tags`). See `nisopaccess(1)`.

`nis_getservlist()` returns a null terminated list of `nis_server` structures that represent the list of servers that serve the domain named `dirname`. Servers from this list can be used when calling functions that require the name of a NIS+ server. For a description of the `nis_server` refer to `nis_objects(3NSL)`.

`nis_freeservlist()` frees the list of servers list of servers returned by `nis_getservlist()`. Note that this is the only legal way to free that list.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `nisopaccess(1)`, `nisstat(1M)`, `nis_names(3NSL)`, `nis_objects(3NSL)`, `nis_subr(3NSL)`, `attributes(5)`

**NOTES** NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <http://www.sun.com/directory/nisplus/transition.html>.

## nis\_subr(3NSL)

<b>NAME</b>	nis_subr, nis_leaf_of, nis_name_of, nis_domain_of, nis_getnames, nis_freenames, nis_dir_cmp, nis_clone_object, nis_destroy_object, nis_print_object – NIS+ subroutines
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpcsvc/nis.h&gt;  nis_name nis_leaf_of(const nis_name name); nis_name nis_name_of(const nis_name name); nis_name nis_domain_of(const nis_name name); nis_name *nis_getnames(const nis_name name); void nis_freenames(nis_name *namelist); name_pos nis_dir_cmp(const nis_name n1, const nis_name n2); nis_object *nis_clone_object(const nis_object *src, nis_object     *dest); void nis_destroy_object(nis_object *obj); void nis_print_object(const nis_object *obj);</pre>
<b>DESCRIPTION</b>	<p>These subroutines are provided to assist in the development of NIS+ applications. They provide several useful operations on both NIS+ names and objects.</p> <p>The first group, <code>nis_leaf_of()</code>, <code>nis_domain_of()</code>, and <code>nis_name_of()</code> provide the functions for parsing NIS+ names. <code>nis_leaf_of()</code> will return the first label in an NIS+ name. It takes into account the double quote character "" which can be used to protect embedded '.' (dot) characters in object names. Note that the name returned will never have a trailing dot character. If passed the global root directory name ".", it will return the null string.</p> <p><code>nis_domain_of()</code> returns the name of the NIS+ domain in which an object resides. This name will always be a fully qualified NIS+ name and ends with a dot. By iteratively calling <code>nis_leaf_of()</code> and <code>nis_domain_of()</code> it is possible to break a NIS+ name into its individual components.</p> <p><code>nis_name_of()</code> is used to extract the unique part of a NIS+ name. This function removes from the tail portion of the name all labels that are in common with the local domain. Thus if a machine were in domain <code>foo.bar.baz.</code> and <code>nis_name_of()</code> were passed a name <code>bob.friends.foo.bar.baz</code>, then <code>nis_name_of()</code> would return the unique part, <code>bob.friends</code>. If the name passed to this function is not in either the local domain or one of its children, this function will return null.</p> <p><code>nis_getnames()</code> will return a list of candidate names for the name passed in as <i>name</i>. If this name is not fully qualified, <code>nis_getnames()</code> will generate a list of names using the default NIS+ directory search path, or the environment variable <code>NIS_PATH</code> if it is set. The returned array of pointers is terminated by a null pointer, and the memory associated with this array should be freed by calling <code>nis_freenames()</code></p>

Though `nis_dir_cmp()` can be used to compare any two NIS+ names, it is used primarily to compare domain names. This comparison is done in a case independent fashion, and the results are an enum of type `name_pos`. When the names passed to this function are identical, the function returns a value of `SAME_NAME`. If the name *n1* is a direct ancestor of name *n2*, then this function returns the result `HIGHER_NAME`. Similarly, if the name *n1* is a direct descendant of name *n2*, then this function returns the result `LOWER_NAME`. When the name *n1* is neither a direct ancestor nor a direct descendant of *n2*, as it would be if the two names were siblings in separate portions of the namespace, then this function returns the result `NOT_SEQUENTIAL`. Finally, if either name cannot be parsed as a legitimate name then this function returns the value `BAD_NAME`.

The second set of functions, consisting of `nis_clone_object()` and `nis_destroy_object()`, are used for manipulating objects. `nis_clone_object()` creates an exact duplicate of the NIS+ object *src*. If the value of *dest* is non-null, it creates the clone of the object into this object structure and allocate the necessary memory for the variable length arrays. If this parameter is null, a pointer to the cloned object is returned. Refer to `nis_objects(3NSL)` for a description of the `nis_object` structure.

`nis_destroy_object()` can be used to destroy an object created by `nis_clone_object()`. This will free up all memory associated with the object and free the pointer passed. If the object was cloned into an array using the *dest* parameter to `nis_clone_object()`, then the object *cannot* be freed with this function. Instead, the function `xdr_free(xdr_nis_object, dest)` must be used.

`nis_print_object()` prints out the contents of a NIS+ object structure on the standard output. Its primary use is for debugging NIS+ programs.

`nis_leaf_of()`, `nis_name_of()` and `nis_clone_object()` return their results as thread-specific data in multithreaded applications.

#### ENVIRONMENT VARIABLES

**NIS\_PATH** This variable overrides the default NIS+ directory search path used by `nis_getnames()`. It contains an ordered list of directories separated by ':' (colon) characters. The '\$' (dollar sign) character is treated specially. Directory names that end in '\$' have the default domain appended to them, and a '\$' by itself is replaced by the list of directories between the default domain and the global root that are at least two levels deep. The default NIS+ directory search path is '\$'.

#### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `nis_names(3NSL)`, `nis_objects(3NSL)`, `nis_tables(3NSL)`, `attributes(5)`

nis\_subr(3NSL)

**NOTES** | NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <http://www.sun.com/directory/nisplus/transition.html>.

<b>NAME</b>	nis_tables, nis_list, nis_add_entry, nis_remove_entry, nis_modify_entry, nis_first_entry, nis_next_entry – NIS+ table functions
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lnsl [ library ... ] #include &lt;rpcsvc/nis.h&gt;  nis_result *nis_list(nis_name name, uint_t flags, int     (*callback)(nis_name table_name, nis_object *object, void     *userdata), void *userdata);  nis_result *nis_add_entry(nis_name table_name, nis_object *object,     uint_t flags);  nis_result *nis_remove_entry(nis_name name, nis_object *object,     uint_t flags);  nis_result *nis_modify_entry(nis_name name, nis_object *object,     uint_t flags);  nis_result *nis_first_entry(nis_name table_name);  nis_result *nis_next_entry(nis_name table_name, netobj *cookie);  void nis_freeresult(nis_result *result);</pre>
<b>DESCRIPTION</b>	<p>Use the NIS+ table functions to search and modify NIS+ tables. <code>nis_list()</code> is used to search a table in the NIS+ namespace. <code>nis_first_entry()</code> and <code>nis_next_entry()</code> are used to enumerate a table one entry at a time. <code>nis_add_entry()</code>, <code>nis_remove_entry()</code>, and <code>nis_modify_entry()</code> are used to change the information stored in a table. <code>nis_freeresult()</code> is used to free the memory associated with the <code>nis_result</code> structure.</p> <p>Entries within a table are named by NIS+ indexed names. An indexed name is a compound name that is composed of a search criteria and a simple NIS+ name that identifies a table object. A search criteria is a series of column names and their associated values enclosed in bracket '[' ]' characters. Indexed names have the following form:</p> <pre>[ colname=value, . . . ], tablename</pre> <p>The list function, <code>nis_list()</code>, takes an indexed name as the value for the <code>name</code> parameter. Here, the <code>tablename</code> should be a fully qualified NIS+ name unless the <code>EXPAND_NAME</code> flag (described below) is set. The second parameter, <code>flags</code>, defines how the function will respond to various conditions. The value for this parameter is created by logically ORing together one or more flags from the following list.</p> <p><b>FOLLOW_LINKS</b>    If the table specified in <code>name</code> resolves to be a LINK type object (see <code>nis_objects(3NSL)</code>), this flag specifies that the client library follow that link and do the search at that object. If this flag is not set and the name resolves to a link, the error <code>NIS_NOTSEARCHABLE</code> will be returned.</p> <p><b>FOLLOW_PATH</b>    This flag specifies that if the entry is not found within this table, the list operation should follow the path specified in the table</p>

## nis\_tables(3NSL)

	<p>object. When used in conjunction with the <code>ALL_RESULTS</code> flag below, it specifies that the path should be followed regardless of the result of the search. When used in conjunction with the <code>FOLLOW_LINKS</code> flag above, named tables in the path that resolve to links will be followed until the table they point to is located. If a table in the path is not reachable because no server that serves it is available, the result of the operation will be either a “soft” success or a “soft” failure to indicate that not all tables in the path could be searched. If a name in the path names is either an invalid or non-existent object then it is silently ignored.</p>
<code>HARD_LOOKUP</code>	<p>This flag specifies that the operation should continue trying to contact a server of the named table until a definitive result is returned (such as <code>NIS_NOTFOUND</code>).</p>
<code>ALL_RESULTS</code>	<p>This flag can only be used in conjunction with <code>FOLLOW_PATH</code> and a callback function. When specified, it forces all of the tables in the path to be searched. If <i>name</i> does not specify a search criteria (imply that all entries are to be returned), then this flag will cause all of the entries in all of the tables in the path to be returned.</p>
<code>NO_CACHE</code>	<p>This flag specifies that the client library should bypass any client object caches and get its information directly from either the master server or a replica server for the named table.</p>
<code>MASTER_ONLY</code>	<p>This flag is even stronger than <code>NO_CACHE</code> in that it specifies that the client library should <i>only</i> get its information from the master server for a particular table. This guarantees that the information will be up to date. However, there may be severe performance penalties associated with contacting the master server directly on large networks. When used in conjunction with the <code>HARD_LOOKUP</code> flag, this will block the list operation until the master server is up and available.</p>
<code>EXPAND_NAME</code>	<p>When specified, the client library will attempt to expand a partially qualified name by calling <code>nis_getnames()</code>, which uses the environment variable <code>NIS_PATH</code>. See <code>nis_local_names(3NSL)</code>.</p>
<code>RETURN_RESULT</code>	<p>This flag is used to specify that a copy of the returning object be returned in the <code>nis_result</code> structure if the operation was successful.</p>

The third parameter to `nis_list()`, *callback*, is an optional pointer to a function that will process the `ENTRY` type objects that are returned from the search. If this pointer is `NULL`, then all entries that match the search criteria are returned in the `nis_result` structure, otherwise this function will be called once for each entry returned. When called, this function should return 0 when additional objects are desired and 1 when it



no longer wishes to see any more objects. The fourth parameter, *userdata*, is simply passed to callback function along with the returned entry object. The client can use this pointer to pass state information or other relevant data that the callback function might need to process the entries.

The `nis_list()` function is not MT-Safe with callbacks.

`nis_add_entry()` will add the NIS+ object to the NIS+ *table\_name*. The *flags* parameter is used to specify the failure semantics for the add operation. The default (*flags* equal 0) is to fail if the entry being added already exists in the table. The `ADD_OVERWRITE` flag may be used to specify that existing object is to be overwritten if it exists, (a modify operation) or added if it does not exist. With the `ADD_OVERWRITE` flag, this function will fail with the error `NIS_PERMISSION` if the existing object does not allow modify privileges to the client.

If the flag `RETURN_RESULT` has been specified, the server will return a copy of the resulting object if the operation was successful.

`nis_remove_entry()` removes the identified entry from the table or a set of entries identified by *table\_name*. If the parameter *object* is non-null, it is presumed to point to a cached copy of the entry. When the removal is attempted, and the object that would be removed is not the same as the cached object pointed to by *object* then the operation will fail with an `NIS_NOTSAMEOBJ` error. If an object is passed with this function, the search criteria in name is optional as it can be constructed from the values within the entry. However, if no object is present, the search criteria must be included in the *name* parameter. If the flags variable is null, and the search criteria does not uniquely identify an entry, the `NIS_NOTUNIQUE` error is returned and the operation is aborted. If the flag parameter `REM_MULTIPLE` is passed, and if remove permission is allowed for each of these objects, then all objects that match the search criteria will be removed. Note that a null search criteria and the `REM_MULTIPLE` flag will remove all entries in a table.

`nis_modify_entry()` modifies an object identified by *name*. The parameter *object* should point to an entry with the `EN_MODIFIED` flag set in each column that contains new information.

The owner, group, and access rights of an entry are modified by placing the modified information into the respective fields of the parameter, *object*: `zo_owner`, `zo_group`, and `zo_access`.

These columns will replace their counterparts in the entry that is stored in the table. The entry passed must have the same number of columns, same type, and valid data in the modified columns for this operation to succeed.

## nis\_tables(3NSL)

If the flags parameter contains the flag `MOD_SAMEOBJ` then the object pointed to by *object* is assumed to be a cached copy of the original object. If the OID of the object passed is different than the OID of the object the server fetches, then the operation fails with the `NIS_NOTSAMEOBJ` error. This can be used to implement a simple read-modify-write protocol which will fail if the object is modified before the client can write the object back.

If the flag `RETURN_RESULT` has been specified, the server will return a copy of the resulting object if the operation was successful.

`nis_first_entry()` fetches entries from a table one at a time. This mode of operation is extremely inefficient and callbacks should be used instead wherever possible. The table containing the entries of interest is identified by *name*. If a search criteria is present in *name* it is ignored. The value of *cookie* within the `nis_result` structure must be copied by the caller into local storage and passed as an argument to `nis_next_entry()`.

`nis_next_entry()` retrieves the “next” entry from a table specified by *table\_name*. The order in which entries are returned is not guaranteed. Further, should an update occur in the table between client calls to `nis_next_entry()` there is no guarantee that an entry that is added or modified will be seen by the client. Should an entry be removed from the table that would have been the “next” entry returned, the error `NIS_CHAINBROKEN` is returned instead.

The path used when the flag `FOLLOW_PATH` is specified, is the one present in the *first* table searched. The path values in tables that are subsequently searched are ignored.

It is legal to call functions that would access the nameservice from within a list callback. However, calling a function that would itself use a callback, or calling `nis_list()` with a callback from within a list callback function is not currently supported.

There are currently no known methods for `nis_first_entry()` and `nis_next_entry()` to get their answers from only the master server.

The `nis_list()` function is not MT-Safe with callbacks. `nis_list()` callbacks are serialized. A call to `nis_list()` with a callback from within `nis_list()` will deadlock. `nis_list()` with a callback cannot be called from an rpc server. See `rpc_svc_calls(3NSL)`. Otherwise, this function is MT-Safe.

## RETURN VALUES

These functions return a pointer to a structure of type `nis_result`:

```
struct nis_result {
    nis_error    status;
    struct {
        uint_t    objects_len;
        nis_object *objects_val;
    } objects;
    netobj    cookie;
    uint32_t    zticks;
    uint32_t    dticks;
```

```

uint32_t  aticks;
uint32_t  cticks;
};

```

The *status* member contains the error status of the the operation. A text message that describes the error can be obtained by calling the function `nis_sperrno()`. See `nis_error(3NSL)`.

The *objects* structure contains two members. *objects\_val* is an array of *nis\_object* structures; *objects\_len* is the number of cells in the array. These objects will be freed by a call to `nis_freeresult()`. See `nis_names(3NSL)`. If you need to keep a copy of one or more objects, they can be copied with the function `nis_clone_object()` and freed with the function `nis_destroy_object()`. See `nis_server(3NSL)`.

The various ticks contain details of where the time, in microseconds, was taken during a request. They can be used to tune one's data organization for faster access and to compare different database implementations.

*zticks*      The time spent in the NIS+ service itself, this count starts when the server receives the request and stops when it sends the reply.

*dticks*      The time spent in the database backend, this time is measured from the time a database call starts, until a result is returned. If the request results in multiple calls to the database, this is the sum of all the time spent in those calls.

*aticks*      The time spent in any "accelerators" or caches. This includes the time required to locate the server needed to resolve the request.

*cticks*      The total time spent in the request, this clock starts when you enter the client library and stops when a result is returned. By subtracting the sum of the other ticks values from this value you can obtain the local overhead of generating a NIS+ request.

Subtracting the value in *dticks* from the value in *zticks* will yield the time spent in the service code itself. Subtracting the sum of the values in *zticks* and *aticks* from the value in *cticks* will yield the time spent in the client library itself. Note: all of the tick times are measured in microseconds.

**ERRORS**      The client library can return a variety of error returns and diagnostics. The more salient ones are documented below.

`NIS_BADATTRIBUTE`      The name of an attribute did not match up with a named column in the table, or the attribute did not have an associated value.

`NIS_BADNAME`      The name passed to the function is not a legal NIS+ name.

`NIS_BADREQUEST`      A problem was detected in the request structure passed to the client library.

## nis\_tables(3NSL)

NIS_CACHEEXPIRED	The entry returned came from an object cache that has <i>expired</i> . This means that the time to live value has gone to zero and the entry may have changed. If the flag NO_CACHE was passed to the lookup function then the lookup function will retry the operation to get an unexpired copy of the object.
NIS_CBERROR	An RPC error occurred on the server while it was calling back to the client. The transaction was aborted at that time and any unsend data was discarded.
NIS_CBRESULTS	Even though the request was successful, all of the entries have been sent to your callback function and are thus not included in this result.
NIS_FOREIGNNS	The name could not be completely resolved. When the name passed to the function would resolve in a namespace that is outside the NIS+ name tree, this error is returned with a NIS+ object of type DIRECTORY. The returned object contains the type of namespace and contact information for a server within that namespace.
NIS_INVALIDOBJ	The object pointed to by <i>object</i> is not a valid NIS+ entry object for the given table. This could occur if it had a mismatched number of columns, or a different data type than the associated column in the table, for example, binary or text.
NIS_LINKNAMEERROR	The name passed resolved to a LINK type object and the contents of the object pointed to an invalid name.
NIS_MODFAIL	The attempted modification failed for some reason.
NIS_NAMEEXISTS	An attempt was made to add a name that already exists. To add the name, first remove the existing name and then add the new name or modify the existing named object.
NIS_NAMEUNREACHABLE	This soft error indicates that a server for the desired directory of the named table object could not be reached. This can occur when there is a network partition or the server has crashed. Attempting the operation again may succeed. See the HARD_LOOKUP flag.
NIS_NOCALLBACK	The server was unable to contact the callback service on your machine. This results in no data being returned.
NIS_NOMEMORY	Generally a fatal result. It means that the service ran out of heap space.

NIS_NOSUCHNAME	This hard error indicates that the named directory of the table object does not exist. This occurs when the server that should be the parent of the server that serves the table, does not know about the directory in which the table resides.
NIS_NOSUCHTABLE	The named table does not exist.
NIS_NOT_ME	A request was made to a server that does not serve the given name. Normally this will not occur, however if you are not using the built in location mechanism for servers, you may see this if your mechanism is broken.
NIS_NOTFOUND	<p>No entries in the table matched the search criteria. If the search criteria was null (return all entries) then this result means that the table is empty and may safely be removed by calling the <code>nis_remove()</code>.</p> <p>If the <code>FOLLOW_PATH</code> flag was set, this error indicates that none of the tables in the path contain entries that match the search criteria.</p>
NIS_NOTMASTER	<p>A change request was made to a server that serves the name, but it is not the master server. This can occur when a directory object changes and it specifies a new master server. Clients that have cached copies of the directory object in the <code>/var/nis/NIS_SHARED_DIRCACHE</code> file will need to have their cache managers restarted to flush this cache. Use <code>nis_cachemgr -i</code>.</p>
NIS_NOTSAMEOBJ	An attempt to remove an object from the namespace was aborted because the object that would have been removed was not the same object that was passed in the request.
NIS_NOTSEARCHABLE	The table name resolved to a NIS+ object that was not searchable.
NIS_PARTIAL	This result is similar to <code>NIS_NOTFOUND</code> except that it means the request succeeded but resolved to zero entries. When this occurs, the server returns a copy of the table object instead of an entry so that the client may then process the path or implement some other local policy.
NIS_RPCERROR	This fatal error indicates the RPC subsystem failed in some way. Generally there will be a <code>syslog(3C)</code> message indicating why the RPC request failed.

nis\_tables(3NSL)

NIS_S_NOTFOUND	The named entry does not exist in the table, however not all tables in the path could be searched, so the entry may exist in one of those tables.
NIS_S_SUCCESS	Even though the request was successful, a table in the search path was not able to be searched, so the result may not be the same as the one you would have received if that table had been accessible.
NIS_SUCCESS	The request was successful.
NIS_SYSTEMERROR	Some form of generic system error occurred while attempting the request. Check the syslog(3C) record for error messages from the server.
NIS_TOOMANYATTRS	The search criteria passed to the server had more attributes than the table had searchable columns.
NIS_TRYAGAIN	The server connected to was too busy to handle your request. add_entry(), remove_entry(), and modify_entry() return this error when the master server is currently updating its internal state. It can be returned to nis_list() when the function specifies a callback and the server does not have the resources to handle callbacks.
NIS_TYPEMISMATCH	An attempt was made to add or modify an entry in a table, and the entry passed was of a different type than the table.

**ENVIRONMENT VARIABLES**

NIS\_PATH      When set, this variable is the search path used by nis\_list() if the flag EXPAND\_NAME is set.

**ATTRIBUTES**

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe with exceptions

**SEE ALSO**

niscat(1), niserror(1), nismatch(1), nis\_cachemgr(1M), nis\_clone\_object(3NSL), n, nis\_destroy\_object(3NSL), nis\_error(3NSL), nis\_getnames(3NSL), nis\_local\_names(3NSL), nis\_names(3NSL), nis\_objects(3NSL), nis\_server(3NSL), rpc\_svc\_calls(3NSL), syslog(3C), attributes(5)

**WARNINGS**

Use the flag HARD\_LOOKUP carefully since it can cause the application to block indefinitely during a network partition.

**NOTES** NIS+ might not be supported in future releases of the Solaris™ Operating Environment. Tools to aid the migration from NIS+ to LDAP are available in the Solaris 9 operating environment. For more information, visit <http://www.sun.com/directory/nisplus/transition.html>.

## nlsgetcall(3NSL)

<b>NAME</b>	nlsgetcall – get client’s data passed via the listener				
<b>SYNOPSIS</b>	<pre>#include &lt;sys/tiuser.h&gt; struct t_call *nlsgetcall(int <i>fildev</i>);</pre>				
<b>DESCRIPTION</b>	<p>nlsgetcall() allows server processes started by the listener process to access the client’s t_call structure, that is, the <i>sndcall</i> argument of t_connect(3NSL).</p> <p>The t_call structure returned by nlsgetcall() can be released using t_free(3NSL).</p> <p>nlsgetcall() returns the address of an allocated t_call structure or NULL if a t_call structure cannot be allocated. If the t_alloc() succeeds, undefined environment variables are indicated by a negative <i>len</i> field in the appropriate netbuf structure. A <i>len</i> field of zero in the netbuf structure is valid and means that the original buffer in the listener’s t_call structure was NULL.</p>				
<b>RETURN VALUES</b>	A NULL pointer is returned if a t_call structure cannot be allocated by t_alloc(). t_errno can be inspected for further error information. Undefined environment variables are indicated by a negative length field ( <i>len</i> ) in the appropriate netbuf structure.				
<b>FILES</b>	<pre>/usr/lib/libnsl_s.a /usr/lib/libslan.a /usr/lib/libnls.a</pre>				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>Unsafe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Unsafe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Unsafe				
<b>SEE ALSO</b>	nlsadmin(1M), getenv(3C), t_alloc(3NSL), t_connect(3NSL), t_error(3NSL), t_free(3NSL), t_sync(3NSL), attributes(5)				
<b>WARNINGS</b>	<p>The <i>len</i> field in the netbuf structure is defined as being unsigned. In order to check for error returns, it should first be cast to an int.</p> <p>The listener process limits the amount of user data (<i>udata</i>) and options data (<i>opt</i>) to 128 bytes each. Address data <i>addr</i> is limited to 64 bytes. If the original data was longer, no indication of overflow is given.</p>				
<b>NOTES</b>	<p>Server processes must call t_sync(3NSL) before calling this routine.</p> <p>This interface is unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.</p>				



<b>NAME</b>	nlsprovider – get name of transport provider				
<b>SYNOPSIS</b>	char *nlsprovider(void);				
<b>DESCRIPTION</b>	<p>nlsprovider() returns a pointer to a null-terminated character string which contains the name of the transport provider as placed in the environment by the listener process. If the variable is not defined in the environment, a NULL pointer is returned.</p> <p>The environment variable is only available to server processes started by the listener process.</p>				
<b>RETURN VALUES</b>	If the variable is not defined in the environment, a NULL pointer is returned.				
<b>FILES</b>	<p>/usr/lib/libslan.a (7300)</p> <p>/usr/lib/libnls.a (3B2 Computer)</p> <p>/usr/lib/libnsl_s.a</p>				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Unsafe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Unsafe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Unsafe				
<b>SEE ALSO</b>	nlsadmin(1M), attributes(5)				
<b>NOTES</b>	This interface is unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.				

## nlsrequest(3NSL)

**NAME** nlsrequest – format and send listener service request message

**SYNOPSIS**

```
#include <listen.h>

int nlsrequest(int fildes, char *service_code);

extern int _nlslogt_errno;
extern char *_nlsrmsg;
```

**DESCRIPTION** Given a virtual circuit to a listener process (*fildes*) and a service code of a server process, `nlsrequest()` formats and sends a *service request message* to the remote listener process requesting that it start the given service. `nlsrequest()` waits for the remote listener process to return a *service request response message*, which is made available to the caller in the static, null-terminated data buffer pointed to by `_nlsrmsg`. The *service request response message* includes a success or failure code and a text message. The entire message is printable.

**RETURN VALUES** The success or failure code is the integer return code from `nlsrequest()`. Zero indicates success, other negative values indicate `nlsrequest()` failures as follows:

- 1 Error encountered by `nlsrequest()`, see `t_errno`.

Positive values are error return codes from the *listener* process. Mnemonics for these codes are defined in `<listen.h>`.

- 2 Request message not interpretable.
- 3 Request service code unknown.
- 4 Service code known, but currently disabled.

If non-null, `_nlsrmsg` contains a pointer to a static, null-terminated character buffer containing the *service request response message*. Note that both `_nlsrmsg` and the data buffer are overwritten by each call to `nlsrequest()`.

If `_nlslog` is non-zero, `nlsrequest()` prints error messages on `stderr`. Initially, `_nlslog` is zero.

**FILES**

- `/usr/lib/libnls.a`
- `/usr/lib/libslan.a`
- `/usr/lib/libnsl_s.a`

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO** `nlsadmin(1M)`, `t_error(3NSL)`, `t_snd(3NSL)`, `t_rcv(3NSL)`, `attributes(5)`

**WARNINGS** | `nlsrequest()` cannot always be certain that the remote server process has been successfully started. In this case, `nlsrequest()` returns with no indication of an error and the caller will receive notification of a disconnect event by way of a `T_LOOK` error before or during the first `t_snd()` or `t_rcv()` call.

**NOTES** | These interfaces are unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.

## rcmd(3SOCKET)

<b>NAME</b>	<code>rcmd</code> , <code>rcmd_af</code> , <code>rresvport</code> , <code>rresvport_af</code> , <code>ruserok</code> – routines for returning a stream to a remote command
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ]  int rcmd(char **ahost, unsigned short inport, const char *luser, const char *ruser, const char *cmd, int *fd2p);  int rcmd_af(char **ahost, unsigned short inport, const char *luser, const char *ruser, const char *cmd, int *fd2p, int af);  int rresvport(int *port);  int rresvport_af(int *port, int af);  int ruserok(const char *rhost, int suser, const char *ruser, const char *luser);</pre>
<b>DESCRIPTION</b>	<p><code>rcmd()</code> is a routine used by the superuser to execute a command on a remote machine using an authentication scheme based on reserved port numbers. It is assumed that an <code>AF_INET</code> socket is returned with <code>rcmd()</code>. <code>rcmd_af()</code> allows the application to choose which type of socket is returned by passing in the address family, either <code>AF_INET</code> or <code>AF_INET6</code>.</p> <p><code>rresvport()</code> is a routine that returns a descriptor to a socket with an address in the privileged port space. <code>rresvport_af()</code> is equivalent to <code>rresvport()</code>, except that you can choose the type of socket address family that will be returned by <code>rresvport_af()</code>, either <code>AF_INET</code> or <code>AF_INET6</code>.</p> <p><code>ruserok()</code> is a routine used by servers to authenticate clients requesting service with <code>rcmd</code>.</p> <p>All of these functions are present in the same file and are used by the <code>in.rshd(1M)</code> server (among others).</p> <p><code>rcmd()</code> and <code>rcmd_af()</code> look up the host <i>ahost</i> using <code>getipnodebyname(3SOCKET)</code>, returning <code>-1</code> if the host does not exist. Otherwise <i>ahost</i> is set to the standard name of the host and a connection is established to a server residing at the well-known Internet port <i>inport</i>.</p> <p>If the connection succeeds, a socket in the Internet domain of type <code>SOCK_STREAM</code> is returned to the caller, and given to the remote command as its standard input (file descriptor 0) and standard output (file descriptor 1). If <i>fd2p</i> is non-zero, then an auxiliary channel to a control process will be set up, and a descriptor for it will be placed in <i>fd2p</i>. The control process will return diagnostic output from the command (file descriptor 2) on this channel, and will also accept bytes on this channel as signal numbers, to be forwarded to the process group of the command. If <i>fd2p</i> is 0, then the standard error (file descriptor 2) of the remote command will be made the same as its standard output and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.</p> <p>The protocol is described in detail in <code>in.rshd(1M)</code>.</p>

The `rresvport()` and `rresvport_af()` routines are used to obtain a socket bound to a privileged port number. This socket is suitable for use by `rcmd()` and `rresvport_af()` and several other routines. Privileged Internet ports are those in the range 1 to 1023. Only the superuser is allowed to bind a socket to a privileged port number. The application must pass in *port*, which must be in the range 512 to 1023. The system first tries to bind to that port number. If it fails, the system then tries to bind to another unused privileged port, if one is available.

`ruserok()` takes a remote host's name, as returned by a `gethostbyaddr()` routine, two user names and a flag indicating whether the local user's name is that of the superuser. See `gethostbyname(3NSL)`. It then checks the files `/etc/hosts.equiv` and possibly `~/.rhosts` in the local user's home directory to see if the request for service is allowed. 0 is returned if the machine name is listed in the `/etc/hosts.equiv` file, or the host and remote user name are found in the `~/.rhosts` file; otherwise `ruserok()` returns -1. If the superuser flag is 1, the checking of the `/etc/hosts.equiv` file is bypassed.

**RETURN VALUES**

`rcmd()` and `rcmd_af()` return a valid socket descriptor upon success. They return -1 upon error and print a diagnostic message to standard error.

`rresvport()` and `rresvport_af()` return a valid, bound socket descriptor upon success. They return -1 upon error with the global value `errno` set according to the reason for failure.

**FILES**

`/etc/hosts.equiv`            system trusted hosts and users  
`~/.rhosts`                    user's trusted hosts and users

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO**

`rlogin(1)`, `rsh(1)`, `in.rexecd(1M)`, `in.rshd(1M)`, `intro(2)`, `gethostbyname(3NSL)`, `getipnodebyname(3SOCKET)`, `rexec(3SOCKET)`, `attributes(5)`

**NOTES**

The error code `EAGAIN` is overloaded to mean "All network ports in use."

These interfaces are unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.

## recv(3SOCKET)

<b>NAME</b>	recv, recvfrom, recvmsg – receive a message from a socket				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt; #include &lt;sys/uio.h&gt;  ssize_t <b>recv</b>(int s, void *buf, size_t len, int flags);  ssize_t <b>recvfrom</b>(int s, void *buf, size_t len, int flags, struct     sockaddr *from, int *fromlen);  ssize_t <b>recvmsg</b>(int s, struct msghdr *msg, int flags);</pre>				
<b>DESCRIPTION</b>	<p>recv(), recvfrom(), and recvmsg() are used to receive messages from another socket. recv() may be used only on a <i>connected</i> socket (see connect(3SOCKET)), while recvfrom() and recvmsg() may be used to receive data on a socket whether it is in a connected state or not. s is a socket created with socket(3SOCKET).</p> <p>If from is not a NULL pointer, the source address of the message is filled in. fromlen is a value-result parameter, initialized to the size of the buffer associated with from, and modified on return to indicate the actual size of the address stored there. The length of the message is returned. If a message is too long to fit in the supplied buffer, excess bytes may be discarded depending on the type of socket the message is received from. See socket(3SOCKET).</p> <p>If no messages are available at the socket, the receive call waits for a message to arrive, unless the socket is nonblocking, in which case -1 is returned with the external variable errno set to EWOULDBLOCK. Seefcntl(2).</p> <p>The select() call may be used to determine when more data arrives.</p> <p>The flags parameter is formed by ORing one or more of the following:</p> <table><tr><td>MSG_OOB</td><td>Read any “out-of-band” data present on the socket rather than the regular “in-band” data.</td></tr><tr><td>MSG_PEEK</td><td>“Peek” at the data present on the socket; the data is returned, but not consumed, so that a subsequent receive operation will see the same data.</td></tr></table> <p>The recvmsg() call uses a msghdr structure to minimize the number of directly supplied parameters. This structure is defined in &lt;sys/socket.h&gt; and includes the following members:</p> <pre>caddr_t      msg_name;          /* optional address */ int          msg_namelen;       /* size of address */ struct iovec *msg_iov;         /* scatter/gather array */ int          msg_iovlen;        /* # elements in msg_iov */ caddr_t      msg_accrightrights; /* access rights sent/received */ int          msg_accrightrightslen;</pre>	MSG_OOB	Read any “out-of-band” data present on the socket rather than the regular “in-band” data.	MSG_PEEK	“Peek” at the data present on the socket; the data is returned, but not consumed, so that a subsequent receive operation will see the same data.
MSG_OOB	Read any “out-of-band” data present on the socket rather than the regular “in-band” data.				
MSG_PEEK	“Peek” at the data present on the socket; the data is returned, but not consumed, so that a subsequent receive operation will see the same data.				

Here `msg_name` and `msg_namelen` specify the destination address if the socket is unconnected; `msg_name` may be given as a NULL pointer if no names are desired or required. The `msg_iov` and `msg_iovlen` describe the scatter-gather locations, as described in `read(2)`. A buffer to receive any access rights sent along with the message is specified in `msg_accrights`, which has length `msg_accrightslen`.

**RETURN VALUES** These calls return the number of bytes received, or `-1` if an error occurred.

**ERRORS** The `recv()`, `recvfrom()`, and `recvmsg()` functions will fail if:

EBADF	<code>s</code> is an invalid file descriptor.
EINVAL	The <code>MSG_OOB</code> flag is set and no out-of-band data is available.
EINTR	The operation was interrupted by delivery of a signal before any data was available to be received.
EIO	An I/O error occurred while reading from or writing to the file system.
ENOMEM	There was insufficient user memory available for the operation to complete.
ENOSR	There were insufficient STREAMS resources available for the operation to complete.
ENOTSOCK	<code>s</code> is not a socket.
ESTALE	A stale NFS file handle exists.
EWOULDBLOCK	The socket is marked non-blocking and the requested operation would block.

The `recv()` and `recvfrom()` functions will fail if:

EINVAL The `len` argument overflowed a `ssize_t`.

The `recvmsg()` function will fail if:

EINVAL	The <code>msg_iovlen</code> member of the <code>msg_hdr</code> structure pointed to by <code>msg</code> is less than or equal to 0, or greater than <code>[IOV_MAX]</code> . See <code>Intro(2)</code> for a definition of <code>[IOV_MAX]</code> .
EINVAL	Either one of the <code>iov_len</code> values in the <code>msg_iov</code> array member of the <code>msg_hdr</code> structure pointed to by <code>msg</code> was negative, or the sum of the <code>iov_len</code> values in the <code>msg_iov</code> array overflowed a <code>ssize_t</code> .

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

recv(3SOCKET)

**SEE ALSO** | `fcntl(2)`, `ioctl(2)`, `read(2)`, `connect(3SOCKET)`, `getsockopt(3SOCKET)`,  
`send(3SOCKET)`, `socket(3HEAD)`, `socket(3SOCKET)`, `attributes(5)`



<b>NAME</b>	recv – receive a message from a connected socket														
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  ssize_t <b>recv</b>(int <i>socket</i>, void *<i>buffer</i>, size_t <i>length</i>, int <i>flags</i>);</pre>														
<b>DESCRIPTION</b>	<p>The <code>recv()</code> function receives a message from a connection-mode or connectionless-mode socket. It is normally used with connected sockets because it does not permit the application to retrieve the source address of received data. The function takes the following arguments:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><i>socket</i></td> <td>Specifies the socket file descriptor.</td> </tr> <tr> <td><i>buffer</i></td> <td>Points to a buffer where the message should be stored.</td> </tr> <tr> <td><i>length</i></td> <td>Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.</td> </tr> <tr> <td><i>flags</i></td> <td>Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:</td> </tr> <tr> <td style="padding-left: 40px;">MSG_PEEK</td> <td>Peeks at an incoming message. The data is treated as unread and the next <code>recv()</code> or similar function will still return this data.</td> </tr> <tr> <td style="padding-left: 40px;">MSG_OOB</td> <td>Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.</td> </tr> <tr> <td style="padding-left: 40px;">MSG_WAITALL</td> <td>Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.</td> </tr> </table> <p>The <code>recv()</code> function returns the length of the message written to the buffer pointed to by the <i>buffer</i> argument. For message-based sockets such as SOCK_DGRAM and SOCK_SEQPACKET, the entire message must be read in a single operation. If a message is too long to fit in the supplied buffer, and MSG_PEEK is not set in the <i>flags</i> argument, the excess bytes are discarded. For stream-based sockets such as SOCK_STREAM, message boundaries are ignored. In this case, data is returned to the user as soon as it becomes available, and no data is discarded.</p> <p>If the MSG_WAITALL flag is not set, data will be returned only up to the end of the first message.</p>	<i>socket</i>	Specifies the socket file descriptor.	<i>buffer</i>	Points to a buffer where the message should be stored.	<i>length</i>	Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.	<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:	MSG_PEEK	Peeks at an incoming message. The data is treated as unread and the next <code>recv()</code> or similar function will still return this data.	MSG_OOB	Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.	MSG_WAITALL	Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.
<i>socket</i>	Specifies the socket file descriptor.														
<i>buffer</i>	Points to a buffer where the message should be stored.														
<i>length</i>	Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.														
<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:														
MSG_PEEK	Peeks at an incoming message. The data is treated as unread and the next <code>recv()</code> or similar function will still return this data.														
MSG_OOB	Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.														
MSG_WAITALL	Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.														

## recv(3XNET)

If no messages are available at the socket and `O_NONBLOCK` is not set on the socket's file descriptor, `recv()` blocks until a message arrives. If no messages are available at the socket and `O_NONBLOCK` is set on the socket's file descriptor, `recv()` fails and sets `errno` to `EAGAIN` or `EWOULDBLOCK`.

**USAGE** The `recv()` function is identical to `recvfrom(3XNET)` with a zero `address_len` argument, and to `read()` if no flags are used.

The `select(3C)` and `poll(2)` functions can be used to determine when data is available to be received.

**RETURN VALUES** Upon successful completion, `recv()` returns the length of the message in bytes. If no messages are available to be received and the peer has performed an orderly shutdown, `recv()` returns 0. Otherwise, `-1` is returned and `errno` is set to indicate the error.

**ERRORS** The `recv()` function will fail if:

`EAGAIN`

`EWOULDBLOCK`

The socket's file descriptor is marked `O_NONBLOCK` and no data is waiting to be received; or `MSG_OOB` is set and no out-of-band data is available and either the socket's file descriptor is marked `O_NONBLOCK` or the socket does not support blocking to await out-of-band data.

`EBADF`

The *socket* argument is not a valid file descriptor.

`ECONNRESET`

A connection was forcibly closed by a peer.

`EFAULT`

The *buffer* parameter can not be accessed or written.

`EINTR`

The `recv()` function was interrupted by a signal that was caught, before any data was available.

`EINVAL`

The `MSG_OOB` flag is set and no out-of-band data is available.

`ENOTCONN`

A receive is attempted on a connection-mode socket that is not connected.

`ENOTSOCK`

The *socket* argument does not refer to a socket.

`EOPNOTSUPP`

The specified flags are not supported for this socket type or protocol.

`ETIMEDOUT`

The connection timed out during connection establishment, or due to a transmission timeout on active connection.

The `recv()` function may fail if:

recv(3XNET)

EIO An I/O error occurred while reading from or writing to the file system.

ENOBUFS Insufficient resources were available in the system to perform the operation.

ENOMEM Insufficient memory was available to fulfill the request.

ENOSR There were insufficient STREAMS resources available for the operation to complete.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `poll(2)`, `recvmsg(3XNET)`, `recvfrom(3XNET)`, `select(3C)`, `send(3XNET)`, `sendmsg(3XNET)`, `sendto(3XNET)`, `shutdown(3XNET)`, `socket(3XNET)`, `attributes(5)`

## recvfrom(3XNET)

<b>NAME</b>	recvfrom – receive a message from a socket												
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  ssize_t <b>recvfrom</b>(int <i>socket</i>, void *<i>buffer</i>, size_t <i>length</i>, int <i>flags</i>,                   struct sockaddr *<i>address</i>, socklen_t *<i>address_len</i>);</pre>												
<b>DESCRIPTION</b>	<p>The <code>recvfrom()</code> function receives a message from a connection-mode or connectionless-mode socket. It is normally used with connectionless-mode sockets because it permits the application to retrieve the source address of received data.</p> <p>The function takes the following arguments:</p> <table><tr><td><i>socket</i></td><td>Specifies the socket file descriptor.</td></tr><tr><td><i>buffer</i></td><td>Points to the buffer where the message should be stored.</td></tr><tr><td><i>length</i></td><td>Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.</td></tr><tr><td><i>flags</i></td><td>Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:  MSG_PEEK                      Peeks at an incoming message. The data is treated as unread and the next <code>recvfrom()</code> or similar function will still return this data.  MSG_OOB                        Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.  MSG_WAITALL                   Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.</td></tr><tr><td><i>address</i></td><td>A null pointer, or points to a <code>sockaddr</code> structure in which the sending address is to be stored. The length and format of the address depend on the address family of the socket.</td></tr><tr><td><i>address_len</i></td><td>Specifies the length of the <code>sockaddr</code> structure pointed to by the <i>address</i> argument.</td></tr></table>	<i>socket</i>	Specifies the socket file descriptor.	<i>buffer</i>	Points to the buffer where the message should be stored.	<i>length</i>	Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.	<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:  MSG_PEEK                      Peeks at an incoming message. The data is treated as unread and the next <code>recvfrom()</code> or similar function will still return this data.  MSG_OOB                        Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.  MSG_WAITALL                   Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.	<i>address</i>	A null pointer, or points to a <code>sockaddr</code> structure in which the sending address is to be stored. The length and format of the address depend on the address family of the socket.	<i>address_len</i>	Specifies the length of the <code>sockaddr</code> structure pointed to by the <i>address</i> argument.
<i>socket</i>	Specifies the socket file descriptor.												
<i>buffer</i>	Points to the buffer where the message should be stored.												
<i>length</i>	Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.												
<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:  MSG_PEEK                      Peeks at an incoming message. The data is treated as unread and the next <code>recvfrom()</code> or similar function will still return this data.  MSG_OOB                        Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.  MSG_WAITALL                   Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.												
<i>address</i>	A null pointer, or points to a <code>sockaddr</code> structure in which the sending address is to be stored. The length and format of the address depend on the address family of the socket.												
<i>address_len</i>	Specifies the length of the <code>sockaddr</code> structure pointed to by the <i>address</i> argument.												

The `recvfrom()` function returns the length of the message written to the buffer pointed to by the *buffer* argument. For message-based sockets such as `SOCK_DGRAM` and `SOCK_SEQPACKET`, the entire message must be read in a single operation. If a message is too long to fit in the supplied buffer, and `MSG_PEEK` is not set in the *flags* argument, the excess bytes are discarded. For stream-based sockets such as `SOCK_STREAM`, message boundaries are ignored. In this case, data is returned to the user as soon as it becomes available, and no data is discarded.

If the `MSG_WAITALL` flag is not set, data will be returned only up to the end of the first message.

Not all protocols provide the source address for messages. If the *address* argument is not a null pointer and the protocol provides the source address of messages, the source address of the received message is stored in the `sockaddr` structure pointed to by the *address* argument, and the length of this address is stored in the object pointed to by the *address\_len* argument.

If the actual length of the address is greater than the length of the supplied `sockaddr` structure, the stored address will be truncated.

If the *address* argument is not a null pointer and the protocol does not provide the source address of messages, the value stored in the object pointed to by *address* is unspecified.

If no messages are available at the socket and `O_NONBLOCK` is not set on the socket's file descriptor, `recvfrom()` blocks until a message arrives. If no messages are available at the socket and `O_NONBLOCK` is set on the socket's file descriptor, `recvfrom()` fails and sets `errno` to `EAGAIN` or `EWOULDBLOCK`.

**USAGE** The `select(3C)` and `poll(2)` functions can be used to determine when data is available to be received.

**RETURN VALUES** Upon successful completion, `recvfrom()` returns the length of the message in bytes. If no messages are available to be received and the peer has performed an orderly shutdown, `recvfrom()` returns 0. Otherwise the function returns `-1` and sets `errno` to indicate the error.

**ERRORS** The `recvfrom()` function will fail if:

`EAGAIN`

`EWOULDBLOCK`

The socket's file descriptor is marked `O_NONBLOCK` and no data is waiting to be received; or `MSG_OOB` is set and no out-of-band data is available and either the socket's file descriptor is marked `O_NONBLOCK` or the socket does not support blocking to await out-of-band data.

`EBADF`

The *socket* argument is not a valid file descriptor.

`ECONNRESET`

A connection was forcibly closed by a peer.

## recvfrom(3XNET)

EFAULT	The <i>buffer</i> , <i>address</i> or <i>address_len</i> parameter can not be accessed or written.
EINTR	A signal interrupted <code>recvfrom()</code> before any data was available.
EINVAL	The MSG_OOB flag is set and no out-of-band data is available.
ENOTCONN	A receive is attempted on a connection-mode socket that is not connected.
ENOTSOCK	The <i>socket</i> argument does not refer to a socket.
EOPNOTSUPP	The specified flags are not supported for this socket type.
ETIMEDOUT	The connection timed out during connection establishment, or due to a transmission timeout on active connection.

The `recvfrom()` function may fail if:

EIO	An I/O error occurred while reading from or writing to the file system.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
ENOMEM	Insufficient memory was available to fulfill the request.
ENOSR	There were insufficient STREAMS resources available for the operation to complete.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `poll(2)`, `recv(3XNET)`, `recvmsg(3XNET)`, `select(3C)`, `send(3XNET)`, `sendmsg(3XNET)`, `sendto(3XNET)`, `shutdown(3XNET)`, `socket(3XNET)`, `attributes(5)`

<b>NAME</b>	recvmsg – receive a message from a socket												
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  ssize_t <b>recvmsg</b>(int <i>socket</i>, struct msghdr *<i>message</i>, int <i>flags</i>);</pre>												
<b>DESCRIPTION</b>	<p>The <code>recvmsg()</code> function receives a message from a connection-mode or connectionless-mode socket. It is normally used with connectionless-mode sockets because it permits the application to retrieve the source address of received data.</p> <p>The function takes the following arguments:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><i>socket</i></td> <td>Specifies the socket file descriptor.</td> </tr> <tr> <td><i>message</i></td> <td>Points to a <code>msghdr</code> structure, containing both the buffer to store the source address and the buffers for the incoming message. The length and format of the address depend on the address family of the socket. The <code>msg_flags</code> member is ignored on input, but may contain meaningful values on output.</td> </tr> <tr> <td><i>flags</i></td> <td>Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:</td> </tr> <tr> <td style="padding-left: 40px;">MSG_OOB</td> <td>Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.</td> </tr> <tr> <td style="padding-left: 40px;">MSG_PEEK</td> <td>Peeks at the incoming message.</td> </tr> <tr> <td style="padding-left: 40px;">MSG_WAITALL</td> <td>Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if <code>MSG_PEEK</code> was specified, or if an error is pending for the socket.</td> </tr> </table> <p>The <code>recvmsg()</code> function receives messages from unconnected or connected sockets and returns the length of the message.</p> <p>The <code>recvmsg()</code> function returns the total length of the message. For message-based sockets such as <code>SOCK_DGRAM</code> and <code>SOCK_SEQPACKET</code>, the entire message must be read in a single operation. If a message is too long to fit in the supplied buffers, and <code>MSG_PEEK</code> is not set in the <i>flags</i> argument, the excess bytes are discarded, and <code>MSG_TRUNC</code> is set in the <code>msg_flags</code> member of the <code>msghdr</code> structure. For stream-based sockets such as <code>SOCK_STREAM</code>, message boundaries are ignored. In this case, data is returned to the user as soon as it becomes available, and no data is discarded.</p>	<i>socket</i>	Specifies the socket file descriptor.	<i>message</i>	Points to a <code>msghdr</code> structure, containing both the buffer to store the source address and the buffers for the incoming message. The length and format of the address depend on the address family of the socket. The <code>msg_flags</code> member is ignored on input, but may contain meaningful values on output.	<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:	MSG_OOB	Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.	MSG_PEEK	Peeks at the incoming message.	MSG_WAITALL	Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if <code>MSG_PEEK</code> was specified, or if an error is pending for the socket.
<i>socket</i>	Specifies the socket file descriptor.												
<i>message</i>	Points to a <code>msghdr</code> structure, containing both the buffer to store the source address and the buffers for the incoming message. The length and format of the address depend on the address family of the socket. The <code>msg_flags</code> member is ignored on input, but may contain meaningful values on output.												
<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:												
MSG_OOB	Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.												
MSG_PEEK	Peeks at the incoming message.												
MSG_WAITALL	Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if <code>MSG_PEEK</code> was specified, or if an error is pending for the socket.												

## recvmsg(3XNET)

If the `MSG_WAITALL` flag is not set, data will be returned only up to the end of the first message.

If no messages are available at the socket, and `O_NONBLOCK` is not set on the socket's file descriptor, `recvmsg()` blocks until a message arrives. If no messages are available at the socket and `O_NONBLOCK` is set on the socket's file descriptor, the `recvmsg()` function fails and sets `errno` to `EAGAIN` or `EWOULDBLOCK`.

In the `msg_hdr` structure, the `msg_name` and `msg_namelen` members specify the source address if the socket is unconnected. If the socket is connected, the `msg_name` and `msg_namelen` members are ignored. The `msg_name` member may be a null pointer if no names are desired or required. The `msg_iov` and `msg_iovlen` fields are used to specify where the received data will be stored. `msg_iov` points to an array of `iovec` structures; `msg_iovlen` must be set to the dimension of this array. In each `iovec` structure, the `iov_base` field specifies a storage area and the `iov_len` field gives its size in bytes. Each storage area indicated by `msg_iov` is filled with received data in turn until all of the received data is stored or all of the areas have been filled.

On successful completion, the `msg_flags` member of the message header is the bitwise-inclusive OR of all of the following flags that indicate conditions detected for the received message:

<code>MSG_EOR</code>	End of record was received (if supported by the protocol).
<code>MSG_OOB</code>	Out-of-band data was received.
<code>MSG_TRUNC</code>	Normal data was truncated.
<code>MSG_CTRUNC</code>	Control data was truncated.

**USAGE** The `select(3C)` and `poll(2)` functions can be used to determine when data is available to be received.

**RETURN VALUES** Upon successful completion, `recvmsg()` returns the length of the message in bytes. If no messages are available to be received and the peer has performed an orderly shutdown, `recvmsg()` returns 0. Otherwise, -1 is returned and `errno` is set to indicate the error.

**ERRORS** The `recvmsg()` function will fail if:

<code>EAGAIN</code> <code>EWOULDBLOCK</code>	The socket's file descriptor is marked <code>O_NONBLOCK</code> and no data is waiting to be received; or <code>MSG_OOB</code> is set and no out-of-band data is available and either the socket's file descriptor is marked <code>O_NONBLOCK</code> or the socket does not support blocking to await out-of-band data.
<code>EBADF</code>	The <i>socket</i> argument is not a valid open file descriptor.
<code>ECONNRESET</code>	A connection was forcibly closed by a peer.



- EFAULT            The *message* parameter, or storage pointed to by the *msg\_name*, *msg\_control* or *msg\_iov* fields of the *message* parameter, or storage pointed to by the *iovec* structures pointed to by the *msg\_iov* field can not be accessed or written.
- EINTR            This function was interrupted by a signal before any data was available.
- EINVAL           The sum of the *iov\_len* values overflows an *ssize\_t*. or the MSG\_OOB flag is set and no out-of-band data is available.
- EMSGSIZE        The *msg\_iovlen* member of the *msghdr* structure pointed to by *message* is less than or equal to 0, or is greater than IOV\_MAX.
- ENOTCONN        A receive is attempted on a connection-mode socket that is not connected.
- ENOTSOCK        The *socket* argument does not refer to a socket.
- EOPNOTSUPP     The specified flags are not supported for this socket type.
- ETIMEDOUT      The connection timed out during connection establishment, or due to a transmission timeout on active connection.

The `recvmsg()` function may fail if:

- EIO              An IO error occurred while reading from or writing to the file system.
- ENOBUFS        Insufficient resources were available in the system to perform the operation.
- ENOMEM         Insufficient memory was available to fulfill the request.
- ENOSR          There were insufficient STREAMS resources available for the operation to complete.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `poll(2)`, `recv(3XNET)`, `recvfrom(3XNET)`, `select(3C)`, `send(3XNET)`, `sendmsg(3XNET)`, `sendto(3XNET)`, `shutdown(3XNET)`, `socket(3XNET)`, `attributes(5)`

resolver(3RESOLV)

<b>NAME</b>	resolver, res_ninit, fp_resstat, res_hostalias, res_nquery, res_nsearch, res_nquerydomain, res_nmkquery, res_nsend, res_nclose, res_nsendsigned, dn_comp, dn_expand, hstrerror, res_init, res_query, res_search, res_mkquery, res_send, herror – resolver routines
<b>BIND 8.2.2 Interfaces</b>	<pre>cc [ flag ... ] file ... -lresolv -lsocket -lnsl [ library ... ] #include &lt;sys/types.h&gt; #include &lt;netinet/in.h&gt; #include &lt;arpa/nameser.h&gt; #include &lt;resolv.h&gt; #include &lt;netdb.h&gt;  int res_ninit(res_state statp);  void fp_resstat(const res_state statp, FILE *fp);  const char *res_hostalias(const res_state statp, const char *name, char * name, char *buf, size_tbuflen);  int res_nquery(res_state statp, const char *dname, int class, int type, u_char *answer, int datalen, int anslen);  int res_nsearch(res_state statp, const char *dname, int class, int type, u_char *answer, int anslen);  int res_nquerydomain(res_state statp, const char *name, const char *domain, int class, int type, u_char *answer, int anslen);  int res_nmkquery(res_state statp, int op, const char *dname, int class, int type, u_char *answer, int datalen, int anslen);  int res_nsend(res_state statp, const u_char *msg, int msglen, u_char *answer, int anslen);  void res_nclose(res_state statp);  int res_nsendsigned(res_state statp, const u_char *msg, int msglen, ns_tsig_key *key, u_char *answer, int anslen);  int dn_comp(const char *exp_dn, u_char *comp_dn, int length, u_char **dnptrs, **lastdnptr);  int dn_expand(const u_char *msg, *eomorig, *comp_dn, char *exp_dn, int length);  const char *hstrerror(int err);</pre>
<b>Deprecated Interfaces</b>	<pre>#include &lt;sys/types.h&gt; #include &lt;netinet/in.h&gt; #include &lt;arpa/nameser.h&gt; #include &lt;resolv.h&gt; #include &lt;netdb.h&gt;  int res_init(void);  int res_query(const char *dname, int class, int type, u_char *answer, int anslen);</pre>

```

int res_search(const char *dname, int class, int type, u_char *answer,
               int anslen);

int res_mkquery(int op, const char *dname, int class, int type, const
                char *data, int datalen, struct rrec *newrr, u_char *buf, int
                buflen);

int res_send(const u_char *msg, int msglen, u_char *answer, int
              anslen);

void herror(const char *s);

```

**DESCRIPTION**

These routines are used for making, sending, and interpreting query and reply messages with Internet domain name servers.

State information is kept in *statp* and is used to control the behavior of these functions. Set *statp* to all zeros prior to making the first call to any of these functions.

The functions `res_init()`, `res_query()`, `res_search()`, `res_mkquery()`, `res_send()`, and `herror()` are deprecated. They are supplied for backwards compatibility. They use global configuration and state information that is kept in the structure `_res` rather than state information referenced through *statp*.

Most of the values in *statp* and `_res` are initialized to reasonable defaults on the first call to `res_ninit()` or `res_init()` and can be ignored. Options stored in `statp->options` or `_res.options` are defined in `<resolv.h>`. They are stored as a simple bit mask containing the bitwise OR of the options enabled.

<code>RES_INIT</code>	True if the initial name server address and default domain name are initialized, that is, <code>res_init()</code> or <code>res_ninit()</code> has been called.
<code>RES_DEBUG</code>	Print debugging messages.
<code>RES_AAONLY</code>	Accept authoritative answers only. With this option, <code>res_send()</code> will continue until it finds an authoritative answer or finds an error. Currently this option is not implemented.
<code>RES_USEVC</code>	Use TCP connections for queries instead of UDP datagrams.
<code>RES_STAYOPEN</code>	Use with <code>RES_USEVC</code> to keep the TCP connection open between queries. This is a useful option for programs that regularly do many queries. The normal mode used should be UDP.
<code>RES_IGNTC</code>	Ignore truncation errors; that is, do not retry with TCP.
<code>RES_RECURSE</code>	Set the recursion-desired bit in queries. This is the default. <code>res_send()</code> and <code>res_nsend()</code> do not do iterative queries and expect the name server to handle recursion.

## resolver(3RESOLV)

<code>RES_DEFNAMES</code>	If set, <code>res_search()</code> and <code>res_nsearch()</code> append the default domain name to single-component names, that is, names that do not contain a dot. This option is enabled by default.
<code>RES_DNSRCH</code>	If this option is set, <code>res_search()</code> and <code>res_nsearch()</code> search for host names in the current domain and in parent domains. See <code>hostname(1)</code> . This option is used by the standard host lookup routine <code>gethostbyname(3NSL)</code> . This option is enabled by default.
<code>RES_NOALIASES</code>	This option turns off the user level aliasing feature controlled by the <code>HOSTALIASES</code> environment variable. Network daemons should set this option.
<code>RES_BLAST</code>	If the <code>RES_BLAST</code> option is defined, <code>resolver()</code> queries will be sent to all servers. If the <code>RES_BLAST</code> option is not defined, but <code>RES_ROTATE</code> is, the list of nameservers are rotated according to a round-robin scheme. <code>RES_BLAST</code> overrides <code>RES_ROTATE</code> .
<code>RES_ROTATE</code>	This option causes <code>res_nsend()</code> and <code>res_send()</code> to rotate the list of nameservers in <code>statp-&gt;nsaddr_list</code> or <code>_res.nsaddr_list</code> .
<code>RES_KEEPTSIG</code>	This option causes <code>res_nsendsigned()</code> to leave the message unchanged after TSIG verification. Otherwise the TSIG record would be removed and the header would be updated.
<b>res_ninit, res_init</b>	The <code>res_ninit()</code> and <code>res_init()</code> routines read the configuration file, if any is present, to get the default domain name, search list and the Internet address of the local name server(s). See <code>resolv.conf(4)</code> . If no server is configured, <code>res_init()</code> or <code>res_ninit()</code> will try to obtain name resolution services from the host on which it is running. The current domain name is defined by <code>domainname(1M)</code> , or by the <code>hostname</code> if it is not specified in the configuration file. Use the environment variable <code>LOCALDOMAIN</code> to override the domain name. This environment variable may contain several blank-separated tokens if you wish to override the search list on a per-process basis. This is similar to the <code>search</code> command in the configuration file. You can set the <code>RES_OPTIONS</code> environment variable to override certain internal resolver options. You can otherwise set them by changing fields in the <code>statp / _res</code> structure. Alternatively, they are inherited from the configuration file's <code>options</code> command. See <code>resolv.conf(4)</code> for information regarding the syntax of the <code>RES_OPTIONS</code> environment variable. Initialization normally occurs on the first call to one of the other resolver routines.
<b>res_nquery, res_query</b>	The <code>res_nquery()</code> and <code>res_query()</code> functions provides interfaces to the server query mechanism. They construct a query, send it to the local server, await a response, and make preliminary checks on the reply. The query requests information of the

specified *type* and *class* for the specified fully-qualified domain name *dname*. The reply message is left in the *answer* buffer with length *anslen* supplied by the caller. `res_nquery()` and `res_query()` return the length of the *answer*, or -1 upon error.

The `res_nquery()` and `res_query()` routines return a length that may be bigger than *anslen*. In that case, retry the query with a larger *buf*. The *answer* to the second query may be larger still], so it is recommended that you supply a *buf* larger than the *answer* returned by the previous query. *answer* must be large enough to receive a maximum UDP response from the server or parts of the *answer* will be silently discarded. The default maximum UDP response size is 512 bytes.

**res\_nsearch,  
res\_search**

The `res_nsearch()` and `res_search()` routines make a query and await a response, just like like `res_nquery()` and `res_query()`. In addition, they implement the default and search rules controlled by the `RES_DEFNAMES` and `RES_DNSRCH` options. They return the length of the first successful reply which is stored in *answer*. On error, they return -1.

The `res_nsearch()` and `res_search()` routines return a length that may be bigger than *anslen*. In that case, retry the query with a larger *buf*. The *answer* to the second query may be larger still], so it is recommended that you supply a *buf* larger than the *answer* returned by the previous query. *answer* must be large enough to receive a maximum UDP response from the server or parts of the *answer* will be silently discarded. The default maximum UDP response size is 512 bytes.

**res\_nmkquery,  
res\_mkquery**

These routines are used by `res_nquery()` and `res_query()`. The `res_nmkquery()` and `res_mkquery()` functions construct a standard query message and place it in *buf*. The routine returns the *size* of the query, or -1 if the query is larger than *buflen*. The query type *op* is usually `QUERY`, but can be any of the query types defined in `<arpa/nameser.h>`. The domain name for the query is given by *dname*. *newrr* is currently unused but is intended for making update messages.

**res\_nsend,  
res\_send,  
res\_nsendsigned**

The `res_nsend()`, `res_send()`, and `res_nsendsigned()` routines send a preformatted query that returns an *answer*. The routine calls `res_ninit()` or `res_init()`. If `RES_INIT` is not set, the routine sends the query to the local name server and handles timeouts and retries. Additionally, the `res_nsendsigned()` uses TSIG signatures to add authentication to the query and verify the response. In this case, only one name server will be contacted. The routines return the length of the reply message, or -1 if there are errors.

The `res_nsend()` and `res_send()` routines return a length that may be bigger than *anslen*. In that case, retry the query with a larger *buf*. The *answer* to the second query may be larger still], so it is recommended that you supply a *buf* larger than the *answer* returned by the previous query. *answer* must be large enough to receive a maximum UDP response from the server or parts of the *answer* will be silently discarded. The default maximum UDP response size is 512 bytes.

**fp\_resstat**

The function `fp_resstat()` prints out the active flag bits in `statp->options` preceded by the text `;; res options:` on *file*.

resolver(3RESOLV)

**res\_hostalias** The function `res_hostalias()` looks up *name* in the file referred to by the HOSTALIASES environment variable and returns the fully qualified host name. If *name* is not found or an error occurs, NULL is returned. `res_hostalias()` stores the result in *buf*.

**res\_nclose** The `res_nclose()` function closes any open files referenced through *statp*.

**dn\_comp** `dn_comp()` compresses the domain name *exp\_dn* and stores it in *comp\_dn*. `dn_comp()` returns the size of the compressed name, or -1 if there were errors. *length* is the size of the array pointed to by *comp\_dn*.

*dnptrs* is a pointer to the head of the list of pointers to previously compressed names in the current message. The first pointer must point to the beginning of the message. The list ends with NULL. The limit to the array is specified by *lastdnptr*.

A side effect of calling `dn_comp()` is to update the list of pointers for labels inserted into the message by `dn_comp()` as the name is compressed. If *dnptrs* is NULL, names are not compressed. If *lastdnptr* is NULL, `dn_comp()` does not update the list of labels.

**dn\_expand** `dn_expand()` expands the compressed domain name *comp\_dn* to a full domain name. The compressed name is contained in a query or reply message. *msg* is a pointer to the beginning of that message. The uncompressed name is placed in the buffer indicated by *exp\_dn*, which is of size *length*. `dn_expand()` returns the size of the compressed name, or -1 if there was an error.

**hstrerror, herror** The variables *statp->res\_h\_errno* and *\_res.res\_h\_errno* and external variable *h\_errno* are set whenever an error occurs during a resolver operation. The following definitions are given in `<netdb.h>`:

```
#define NETDB_INTERNAL -1 /* see errno */
#define NETDB_SUCCESS 0 /* no problem */
#define HOST_NOT_FOUND 1 /* Authoritative Answer Host not found */
#define TRY_AGAIN 2 /* Non-Authoritative not found, or SERVFAIL */
#define NO_RECOVERY 3 /* Non-Recoverable: FORMERR, REFUSED, NOTIMP*/
#define NO_DATA 4 /* Valid name, no data for requested type */
```

The `herror()` function writes a message to the diagnostic output consisting of the string parameters, the constant string ":", and a message corresponding to the value of *h\_errno*.

The `hstrerror()` function returns a string, which is the message text that corresponds to the value of the *err* parameter.

**FILES** /etc/resolv.conf

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsl (32-bit)

resolver(3RESOLV)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
	SUNWcslx (64-bit)
Interface Stability	Standard BIND 8.2.4
MT-Level	Unsafe for Deprecated Interfaces; MT-Safe for all others.

**SEE ALSO** domainname(1M), in.named(1M), gethostbyname(3NSL), libresolv(3LIB), resolv.conf(4), attributes(5)

Lottor, M. *RFC 1033, Domain Administrators Operations Guide*. Network Working Group. November 1987.

Mockapetris, Paul. *RFC 1034, Domain Names - Concepts and Facilities*. Network Working Group. November 1987.

Mockapetris, Paul. *RFC 1035, Domain Names - Implementation and Specification*. Network Working Group. November 1987.

Partridge, Craig. *RFC 974, Mail Routing and the Domain System*. Network Working Group. January 1986.

Stahl, M. *RFC 1032, Domain Administrators Guide*. Network Working Group. November 1987.

Vixie, Paul, Dunlap, Kevin J., Karels, Michael J. *Name Server Operations Guide for BIND*. Internet Software Consortium, 1996.

**NOTES** When the caller supplies a work buffer, for example the *answer* buffer argument to `res_nsend()` or `res_send()`, the buffer should be aligned on an eight byte boundary. Otherwise, an error such as a SIGBUS may result.

## rexec(3SOCKET)

<b>NAME</b>	rexec, rexec_af – return stream to a remote command
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ]  int rexec(char **ahost, unsigned short inport, const char *user, const char *passwd, const char *cmd, int *fd2p);  int rexec_af(char **ahost, unsigned short inport, const char *user, const char *passwd, const char *cmd, int *fd2p, int af);</pre>
<b>DESCRIPTION</b>	<p>rexec() and rexec_af() look up the host <i>ahost</i> using getipnodebyname(3SOCKET), returning -1 if the host does not exist. Otherwise <i>ahost</i> is set to the standard name of the host. If a username and password are both specified, then these are used to authenticate to the foreign host; otherwise the user's .netrc file in his home directory is searched for appropriate information. If all this fails, the user is prompted for the information.</p> <p>The difference between rexec() and rexec_af() is that while rexec() always returns a socket of the AF_INET address family, with rexec_af() the application can choose which type of address family the socket returned should be. rexec_af() supports both AF_INET and AF_INET6 address families.</p> <p>The port <i>inport</i> specifies which well-known DARPA Internet port to use for the connection. The port number used must be in network byte order, as supplied by a call to htons(3XNET). The protocol for connection is described in detail in in.rexecd(1M).</p> <p>If the call succeeds, a socket of type SOCK_STREAM is returned to the caller, and given to the remote command as its standard input and standard output. If <i>fd2p</i> is non-zero, then an auxiliary channel will be setup, and a file descriptor for it will be placed in <i>*fd2p</i>. The control process will return diagnostic output (file descriptor 2, the standard error) from the command on this channel, and will also accept bytes on this channel as signal numbers, to be forwarded to the process group of the command. If <i>fd2p</i> is 0, then the standard error (file descriptor 2 of the remote command) will be made the same as its standard output and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.</p>
<b>RETURN VALUES</b>	<p>If rexec() succeeds, a file descriptor number, which is a socket of type SOCK_STREAM and address family AF_INET is returned by the routine. <i>*ahost</i> is set to the standard name of the host, and if <i>fd2p</i> is not NULL, a file descriptor number is placed in <i>*fd2p</i> which represents the command's standard error stream.</p> <p>If rexec_af() succeeds, the routine returns a file descriptor number, which is a socket of type SOCK_STREAM and of address family type AF_INET or AF_INET6, as determined by the value of the <i>af</i> parameter that the caller passes in.</p> <p>If either rexec() or rexec_af() fails, -1 is returned.</p>
<b>ATTRIBUTES</b>	See attributes (5) for descriptions of the following attributes:



rexec(3SOCKET)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO** `in.rexecd(1M)`, `gethostbyname(3NSL)`, `getipnodebyname(3SOCKET)`, `getservbyname(3SOCKET)`, `htons(3XNET)`, `socket(3SOCKET)`, `attributes(5)`

**NOTES** There is no way to specify options to the `socket()` call that `rexec()` or `rexec_af()` makes.

This interface is unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.

## rpc(3NSL)

<b>NAME</b>	rpc – library routines for remote procedure calls						
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lnsl [ <i>library</i> ... ] #include &lt;rpc/rpc.h&gt; #include &lt;netconfig.h&gt;</pre>						
<b>DESCRIPTION</b>	<p>These routines allow C language programs to make procedure calls on other machines across a network. First, the client sends a request to the server. On receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply.</p> <p>All RPC routines require the header <code>&lt;rpc/rpc.h&gt;</code>. Routines that take a <code>netconfig</code> structure also require that <code>&lt;netconfig.h&gt;</code> be included. Applications using RPC and XDR routines should be linked with the <code>libnsl</code> library.</p>						
<b>Multithread Considerations</b>	<p>In the case of multithreaded applications, the <code>-mt</code> option must be specified on the command line at compilation time to enable a thread-specific version of <code>rpc_createerr()</code>. See <code>rpc_clnt_create(3NSL)</code> and <code>threads(3THR)</code>.</p> <p>When used in multithreaded applications, client-side routines are MT-Safe. CLIENT handles can be shared between threads; however, in this implementation, requests by different threads are serialized (that is, the first request will receive its results before the second request is sent). See <code>rpc_clnt_create(3NSL)</code>.</p> <p>When used in multithreaded applications, server-side routines are usually Unsafe. In this implementation the service transport handle, <code>SVCXPRT</code> contains a single data area for decoding arguments and encoding results. See <code>rpc_svc_create(3NSL)</code>. Therefore, this structure cannot be freely shared between threads that call functions that do this. Routines that are affected by this restriction are marked as unsafe for MT applications. See <code>rpc_svc_calls(3NSL)</code>.</p>						
<b>Nettyp</b>	<p>Some of the high-level RPC interface routines take a <i>nettype</i> string as one of the parameters (for example, <code>clnt_create()</code>, <code>svc_create()</code>, <code>rpc_reg()</code>, <code>rpc_call()</code>). This string defines a class of transports which can be used for a particular application.</p> <p><i>nettype</i> can be one of the following:</p> <table><tr><td><code>netpath</code></td><td>Choose from the transports which have been indicated by their token names in the <code>NETPATH</code> environment variable. If <code>NETPATH</code> is unset or <code>NULL</code>, it defaults to <code>visible</code>. <code>netpath</code> is the default <i>nettype</i>.</td></tr><tr><td><code>visible</code></td><td>Choose the transports which have the visible flag (<code>v</code>) set in the <code>/etc/netconfig</code> file.</td></tr><tr><td><code>circuit_v</code></td><td>This is same as <code>visible</code> except that it chooses only the connection oriented transports (semantics <code>tpi_cots</code> or <code>tpi_cots_ord</code>) from the entries in the <code>/etc/netconfig</code> file.</td></tr></table>	<code>netpath</code>	Choose from the transports which have been indicated by their token names in the <code>NETPATH</code> environment variable. If <code>NETPATH</code> is unset or <code>NULL</code> , it defaults to <code>visible</code> . <code>netpath</code> is the default <i>nettype</i> .	<code>visible</code>	Choose the transports which have the visible flag ( <code>v</code> ) set in the <code>/etc/netconfig</code> file.	<code>circuit_v</code>	This is same as <code>visible</code> except that it chooses only the connection oriented transports (semantics <code>tpi_cots</code> or <code>tpi_cots_ord</code> ) from the entries in the <code>/etc/netconfig</code> file.
<code>netpath</code>	Choose from the transports which have been indicated by their token names in the <code>NETPATH</code> environment variable. If <code>NETPATH</code> is unset or <code>NULL</code> , it defaults to <code>visible</code> . <code>netpath</code> is the default <i>nettype</i> .						
<code>visible</code>	Choose the transports which have the visible flag ( <code>v</code> ) set in the <code>/etc/netconfig</code> file.						
<code>circuit_v</code>	This is same as <code>visible</code> except that it chooses only the connection oriented transports (semantics <code>tpi_cots</code> or <code>tpi_cots_ord</code> ) from the entries in the <code>/etc/netconfig</code> file.						

<code>datagram_v</code>	This is same as <code>visible</code> except that it chooses only the connectionless datagram transports (semantics <code>tpi_clts</code> ) from the entries in the <code>/etc/netconfig</code> file.
<code>circuit_n</code>	This is same as <code>netpath</code> except that it chooses only the connection oriented datagram transports (semantics <code>tpi_cots</code> or <code>tpi_cots_ord</code> ).
<code>datagram_n</code>	This is same as <code>netpath</code> except that it chooses only the connectionless datagram transports (semantics <code>tpi_clts</code> ).
<code>udp</code>	This refers to Internet UDP.
<code>tcp</code>	This refers to Internet TCP.

If `nettype` is NULL, it defaults to `netpath`. The transports are tried in left to right order in the `NETPATH` variable or in top to down order in the `/etc/netconfig` file.

### Derived Types

In a 64-bit environment, the derived types are defined as follows:

```
typedef          uint32_t          rpcprog_t;
typedef          uint32_t          rpcvers_t;
typedef          uint32_t          rpcproc_t;
typedef          uint32_t          rpcprot_t;
typedef          uint32_t          rpcport_t;
typedef          int32_t           rpc_inline_t;
```

In a 32-bit environment, the derived types are defined as follows:

```
typedef          unsigned long     rpcprog_t;
typedef          unsigned long     rpcvers_t;
typedef          unsigned long     rpcproc_t;
typedef          unsigned long     rpcprot_t;
typedef          unsigned long     rpcport_t;
typedef          long              rpc_inline_t;
```

### Data Structures

Some of the data structures used by the RPC package are shown below.

#### The AUTH Structure

```
union des_block {
    struct {
        u_int32 high;
        u_int32 low;
```

## rpc(3NSL)

```

        } key;
char c[8];
};
typedef union des_block des_block;
extern bool_t xdr_des_block();
/*
 * Authentication info. Opaque to client.
 */
struct opaque_auth {
    enum_t oa_flavor;        /* flavor of auth */
    caddr_t oa_base;        /* address of more auth stuff */
    uint_t oa_length;       /* not to exceed MAX_AUTH_BYTES */
};
/*
 * Auth handle, interface to client side authenticators.
 */
typedef struct {
    struct opaque_auth ah_cred;
    struct opaque_auth ah_verf;
    union des_block ah_key;
    struct auth_ops {
        void(*ah_nextverf)();
        int(*ah_marshall)();    /* nextverf & serialize */
        int(*ah_validate)();    /* validate verifier */
        int(*ah_refresh)();     /* refresh credentials */
        void(*ah_destroy)();    /* destroy this structure */
    } *ah_ops;
    caddr_t ah_private;
} AUTH;

/*
 * Client rpc handle.
 * Created by individual implementations.
 * Client is responsible for initializing auth.
 */
typedef struct {
    AUTH *cl_auth;        /* authenticator */
    struct clnt_ops {
        enum clnt_stat (*cl_call)();    /* call remote procedure */
        void (*cl_abort)();            /* abort a call */
        void (*cl_geterr)();           /* get specific error code */
        bool_t (*cl_freeres)();        /* frees results */
        void (*cl_destroy)();          /* destroy this structure */
        bool_t (*cl_control)();        /* the ioctl() of rpc */
        int (*cl_settimers)();        /* set rpc level timers */
    } *cl_ops;
    caddr_t cl_private;                /* private stuff */
    char *cl_netid;                    /* network identifier */
    char *cl_tp;                       /* device name */
} CLIENT;

enum xprt_stat {
    XPRT_DIED,
    XPRT_MOREREQS,
    XPRT_IDLE
};
/*
 * Server side transport handle

```

**The CLIENT  
Structure**

**The SVCXPRT  
Structure**

```

*/
typedef struct {
    int xp_fd; /* file descriptor for the
    ushort_t xp_port; /* obsolete */
    struct xp_ops {
        bool_t (*xp_rcv)(); /* receive incoming requests */
        enum xpstat (*xp_stat)(); /* get transport status */
        bool_t (*xp_getargs)(); /* get arguments */
        bool_t (*xp_reply)(); /* send reply */
        bool_t (*xp_freeargs)(); /* free mem allocated
        for args */
        void (*xp_destroy)(); /* destroy this struct */
    } *xp_ops;
    int xp_addrlen; /* length of remote addr.
    Obsolete */
    char *xp_tp; /* transport provider device
    name */
    char *xp_netid; /* network identifier */
    struct netbuf xp_ltaddr; /* local transport address */
    struct netbuf xp_rtaddr; /* remote transport address */
    char xp_raddr[16]; /* remote address. Obsolete */
    struct opaque_auth xp_verf; /* raw response verifier */
    caddr_t xp_p1; /* private: for use
    by svc ops */
    caddr_t xp_p2; /* private: for use
    by svc ops */
    caddr_t xp_p3; /* private: for use
    by svc lib */
    int xp_type /* transport type */
} SVCXPRT;

The svc_req Structure
struct svc_req {
    rpcprog_t rq_prog; /* service program number */
    rpcvers_t rq_vers; /* service protocol version */
    rpcproc_t rq_proc; /* the desired procedure */
    struct opaque_auth rq_cred; /* raw creds from the wire */
    caddr_t rq_clntcred; /* read only cooked cred */
    SVCXPRT *rq_xprt; /* associated transport */
};

The XDR Structure
/*
 * XDR operations.
 * XDR_ENCODE causes the type to be encoded into the stream.
 * XDR_DECODE causes the type to be extracted from the stream.
 * XDR_FREE can be used to release the space allocated by an XDR_DECODE
 * request.
 */
enum xdr_op {
    XDR_ENCODE=0,
    XDR_DECODE=1,
    XDR_FREE=2
};
/*
 * This is the number of bytes per unit of external data.
 */
#define BYTES_PER_XDR_UNIT (4)
#define RNDUP(x) (((x) + BYTES_PER_XDR_UNIT - 1) /

```

## rpc(3NSL)

```

                                BYTES_PER_XDR_UNIT) \ * BYTES_PER_XDR_UNIT)
/*
 * A xdrproc_t exists for each data type which is to be encoded or
 * decoded. The second argument to the xdrproc_t is a pointer to
 * an opaque pointer. The opaque pointer generally points to a
 * structure of the data type to be decoded. If this points to 0,
 * then the type routines should allocate dynamic storage of the
 * appropriate size and return it.
 * bool_t (*xdrproc_t)(XDR *, caddr_t *);
 */
typedef bool_t (*xdrproc_t)();
/*
 * The XDR handle.
 * Contains operation which is being applied to the stream,
 * an operations vector for the particular implementation
 */
typedef struct {
enum xdr_op x_op; /* operation; fast additional param */
struct xdr_ops {

bool_t (*x_getlong)(); /* get long from underlying stream */
bool_t (*x_putlong)(); /* put long to underlying stream */
bool_t (*x_getbytes)(); /* get bytes from underlying stream */
bool_t (*x_putbytes)(); /* put bytes to underlying stream */
uint_t (*x_getpostn)(); /* returns bytes off from beginning */
bool_t (*x_setpostn)(); /* reposition the stream */
rpc_inline_t (*x_inline)(); /* buf quick ptr to buffered data */
void (*x_destroy)(); /* free privates of this xdr_stream */
bool_t (*x_control)(); /* changed/retrieve client object info*/
bool_t (*x_getint32)(); /* get int from underlying stream */
bool_t (*x_putint32)(); /* put int to underlying stream */

} *x_ops;

caddr_t x_public; /* users' data */
caddr_t x_priv /* pointer to private data */
caddr_t x_base; /* private used for position info */
int x_handy; /* extra private word */
XDR;

```

### Index to Routines

The following table lists RPC routines and the manual reference pages on which they are described:

RPC Routine	Manual Reference Page
auth_destroy	rpc_clnt_auth(3NSL)
authdes_create	rpc_soc(3NSL)
authdes_getucred	secure_rpc(3NSL)
authdes_seccreate	secure_rpc(3NSL)
authnone_create	rpc_clnt_auth(3NSL)
authsys_create	rpc_clnt_auth(3NSL)
authsys_create_default	rpc_clnt_auth(3NSL)

authunix_create	rpc_soc(3NSL)
authunix_create_default	rpc_soc(3NSL)
callrpc	rpc_soc(3NSL)
clnt_broadcast	rpc_soc(3NSL)
clnt_call	rpc_clnt_calls(3NSL)
clnt_control	rpc_clnt_create(3NSL)
clnt_create	rpc_clnt_create(3NSL)
clnt_destroy	rpc_clnt_create(3NSL)
clnt_dg_create	rpc_clnt_create(3NSL)
clnt_freeres	rpc_clnt_calls(3NSL)
clnt_geterr	rpc_clnt_calls(3NSL)
clnt_pcreateerror	rpc_clnt_create(3NSL)
clnt_perrno	rpc_clnt_calls(3NSL)
clnt_perror	rpc_clnt_calls(3NSL)
clnt_raw_create	rpc_clnt_create(3NSL)
clnt_spcreateerror	rpc_clnt_create(3NSL)
clnt_sperrno	rpc_clnt_calls(3NSL)
clnt_sperror	rpc_clnt_calls(3NSL)
clnt_tli_create	rpc_clnt_create(3NSL)
clnt_tp_create	rpc_clnt_create(3NSL)
clnt_udpcreate	rpc_soc(3NSL)
clnt_vc_create	rpc_clnt_create(3NSL)
clntraw_create	rpc_soc(3NSL)
clnttcp_create	rpc_soc(3NSL)
clntudp_bufcreate	rpc_soc(3NSL)
get_myaddress	rpc_soc(3NSL)
getnetname	secure_rpc(3NSL)
host2netname	secure_rpc(3NSL)
key_decryptsession	secure_rpc(3NSL)
key_encryptsession	secure_rpc(3NSL)
key_gendes	secure_rpc(3NSL)

## rpc(3NSL)

key_setsecret	secure_rpc(3NSL)
netname2host	secure_rpc(3NSL)
netname2user	secure_rpc(3NSL)
pmap_getmaps	rpc_soc(3NSL)
pmap_getport	rpc_soc(3NSL)
pmap_rmtcall	rpc_soc(3NSL)
pmap_set	rpc_soc(3NSL)
pmap_unset	rpc_soc(3NSL)
rac_drop	rpc_rac(3RAC)
rac_poll	rpc_rac(3RAC)
rac_recv	rpc_rac(3RAC)
rac_send	rpc_rac(3RAC)
registerrpc	rpc_soc(3NSL)
rpc_broadcast	rpc_clnt_calls(3NSL)
rpc_broadcast_exp	rpc_clnt_calls(3NSL)
rpc_call	rpc_clnt_calls(3NSL)
rpc_reg	rpc_svc_calls(3NSL)
svc_create	rpc_svc_create(3NSL)
svc_destroy	rpc_svc_create(3NSL)
svc_dg_create	rpc_svc_create(3NSL)
svc_dg_enablecache	rpc_svc_calls(3NSL)
svc_fd_create	rpc_svc_create(3NSL)
svc_fds	rpc_soc(3NSL)
svc_freeargs	rpc_svc_reg(3NSL)
svc_getargs	rpc_svc_reg(3NSL)
svc_getcaller	rpc_soc(3NSL)
svc_getreq	rpc_soc(3NSL)
svc_getreqset	rpc_svc_calls(3NSL)
svc_getrpccaller	rpc_svc_calls(3NSL)
svc_raw_create	rpc_svc_create(3NSL)
svc_reg	rpc_svc_calls(3NSL)



svc_register	rpc_soc(3NSL)
svc_run	rpc_svc_reg(3NSL)
svc_sendreply	rpc_svc_reg(3NSL)
svc_tli_create	rpc_svc_create(3NSL)
svc_tp_create	rpc_svc_create(3NSL)
svc_unreg	rpc_svc_calls(3NSL)
svc_unregister	rpc_soc(3NSL)
svc_vc_create	rpc_svc_create(3NSL)
svcerr_auth	rpc_svc_err(3NSL)
svcerr_decode	rpc_svc_err(3NSL)
svcerr_noproc	rpc_svc_err(3NSL)
svcerr_noprogram	rpc_svc_err(3NSL)
svcerr_progvers	rpc_svc_err(3NSL)
svcerr_systemerr	rpc_svc_err(3NSL)
svcerr_weakauth	rpc_svc_err(3NSL)
svctcp_create	rpc_soc(3NSL)
svcrow_create	rpc_soc(3NSL)
svctcp_create	rpc_soc(3NSL)
svcudp_bufcreate	rpc_soc(3NSL)
svcudp_create	rpc_soc(3NSL)
user2netname	secure_rpc(3NSL)
xdr_accepted_reply	rpc_xdr(3NSL)
xdr_authsys_parms	rpc_xdr(3NSL)
xdr_authunix_parms	rpc_soc(3NSL)
xdr_callhdr	rpc_xdr(3NSL)
xdr_callmsg	rpc_xdr(3NSL)
xdr_opaque_auth	rpc_xdr(3NSL)
xdr_rejected_reply	rpc_xdr(3NSL)
xdr_replymsg	rpc_xdr(3NSL)
xprt_register	rpc_svc_calls(3NSL)
xprt_unregister	rpc_svc_calls(3NSL)

rpc(3NSL)

**FILES** /etc/netconfig

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe with exceptions

**SEE ALSO** `getnetconfig(3NSL)`, `getnetpath(3NSL)`, `rpc_clnt_auth(3NSL)`, `rpc_clnt_calls(3NSL)`, `rpc_clnt_create(3NSL)`, `rpc_svc_calls(3NSL)`, `rpc_svc_create(3NSL)`, `rpc_svc_err(3NSL)`, `rpc_svc_reg(3NSL)`, `rpc_xdr(3NSL)`, `rpcbind(3NSL)`, `secure_rpc(3NSL)`, `threads(3THR)`, `xdr(3NSL)`, `netconfig(4)`, `rpc(4)`, `attributes(5)`, `environ(5)`

<b>NAME</b>	rpcbind, rpcb_getmaps, rpcb_getaddr, rpcb_gettime, rpcb_rmtcall, rpcb_set, rpcb_unset – library routines for RPC bind service
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpc.h&gt;  struct rpcblist *<b>rpcb_getmaps</b>(const struct netconfig *netconf,     const char *host);  bool_t <b>rpcb_getaddr</b>(const rpcprog_t prognum, const rpcvers_t     versnum, const struct netconfig *netconf, struct netbuf *svcaddr,     const char *host);  bool_t <b>rpcb_gettime</b>(const char *host, time_t *timep);  enum clnt_stat <b>rpcb_rmtcall</b>(const struct netconfig *netconf, const     char *host, const rpcprog_t prognum, const rpcvers_t versnum,     const rpcproc_t procnum, const xdrproc_t inproc, const caddr_t     in, const xdrproc_t outproc, caddr_t out, const struct timeval     tout, struct netbuf *svcaddr);  bool_t <b>rpcb_set</b>(const rpcprog_t prognum, const rpcvers_t versnum,     const struct netconfig *netconf, const struct netbuf *svcaddr);  bool_t <b>rpcb_unset</b>(const rpcprog_t prognum, const rpcvers_t versnum,     const struct netconfig *netconf);</pre>
<b>DESCRIPTION</b>	These routines allow client C programs to make procedure calls to the RPC binder service. rpcbind maintains a list of mappings between programs and their universal addresses. See rpcbind(1M).
<b>Routines</b>	<p><b>rpcb_getmaps()</b> An interface to the rpcbind service, which returns a list of the current RPC program-to-address mappings on <i>host</i>. It uses the transport specified through <i>netconf</i> to contact the remote rpcbind service on <i>host</i>. This routine will return NULL, if the remote rpcbind could not be contacted.</p> <p><b>rpcb_getaddr()</b> An interface to the rpcbind service, which finds the address of the service on <i>host</i> that is registered with program number <i>prognum</i>, version <i>versnum</i>, and speaks the transport protocol associated with <i>netconf</i>. The address found is returned in <i>svcaddr</i>. <i>svcaddr</i> should be preallocated. This routine returns TRUE if it succeeds. A return value of FALSE means that the mapping does not exist or that the RPC system failed to contact the remote rpcbind service. In the latter case, the global variable <i>rpc_createerr</i> contains the RPC status. See <i>rpc_clnt_create</i>(3NSL).</p> <p><b>rpcb_gettime()</b> This routine returns the time on <i>host</i> in <i>timep</i>. If <i>host</i> is NULL, <i>rpcb_gettime()</i> returns the time on its own machine. This routine returns TRUE if it succeeds, FALSE if it fails. <i>rpcb_gettime()</i> can be used to synchronize the time between the client and the remote server. This routine is particularly useful for secure RPC.</p>

## rpcbind(3NSL)

### rpcb\_rmtcall()

An interface to the `rpcbind` service, which instructs `rpcbind` on *host* to make an RPC call on your behalf to a procedure on that host. The `netconfig` structure should correspond to a connectionless transport. The parameter `*svcaddr` will be modified to the server's address if the procedure succeeds. See `rpc_call()` and `clnt_call()` in `rpc_clnt_calls(3NSL)` for the definitions of other parameters.

This procedure should normally be used for a "ping" and nothing else. This routine allows programs to do lookup and call, all in one step.

Note: Even if the server is not running `rpcbind` does not return any error messages to the caller. In such a case, the caller times out.

Note: `rpcb_rmtcall()` is only available for connectionless transports.

### rpcb\_set()

An interface to the `rpcbind` service, which establishes a mapping between the triple [`prognum`, `versnum`, `netconf->nc_netid`] and `svcaddr` on the machine's `rpcbind` service. The value of `nc_netid` must correspond to a network identifier that is defined by the `netconfig` database. This routine returns `TRUE` if it succeeds, `FALSE` otherwise. See also `svc_reg()` in `rpc_svc_calls(3NSL)`. If there already exists such an entry with `rpcbind`, `rpcb_set()` will fail.

### rpcb\_unset()

An interface to the `rpcbind` service, which destroys the mapping between the triple [`prognum`, `versnum`, `netconf->nc_netid`] and the address on the machine's `rpcbind` service. If `netconf` is `NULL`, `rpcb_unset()` destroys all mapping between the triple [`prognum`, `versnum`, `all-transport`] and the addresses on the machine's `rpcbind` service. This routine returns `TRUE` if it succeeds, `FALSE` otherwise. Only the owner of the service or the super-user can destroy the mapping. See also `svc_unreg()` in `rpc_svc_calls(3NSL)`.

## ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

## SEE ALSO

`rpcbind(1M)`, `rpcinfo(1M)`, `rpc_clnt_calls(3NSL)`, `rpc_clnt_create(3NSL)`, `rpc_svc_calls(3NSL)`, `attributes(5)`

<b>NAME</b>	rpc_clnt_auth, auth_destroy, authnone_create, authsys_create, authsys_create_default – library routines for client side remote procedure call authentication
<b>SYNOPSIS</b>	<pre>void <b>auth_destroy</b>(AUTH *auth); AUTH *<b>authnone_create</b>(void); AUTH *<b>authsys_create</b>(const char *host, const uid_t uid, const gid_t     gid, const int len, const gid_t *aup_gids); AUTH *<b>authsys_create_default</b>(void);</pre>
<b>DESCRIPTION</b>	<p>These routines are part of the RPC library that allows C language programs to make procedure calls on other machines across the network, with desired authentication.</p> <p>These routines are normally called after creating the CLIENT handle. The <code>cl_auth</code> field of the CLIENT structure should be initialized by the AUTH structure returned by some of the following routines. The client's authentication information is passed to the server when the RPC call is made.</p> <p>Only the NULL and the SYS style of authentication is discussed here. For the DES style authentication, please refer to <code>secure_rpc(3NSL)</code>.</p> <p>The NULL and SYS style of authentication are safe in multithreaded applications. For the MT-level of the DES style, see its pages.</p>
<b>Routines</b>	<p>The following routines require that the header <code>&lt;rpc/rpc.h&gt;</code> be included (see <code>rpc(3NSL)</code> for the definition of the AUTH data structure).</p> <pre>#include &lt;rpc/rpc.h&gt; <b>auth_destroy</b>()     A function macro that destroys the authentication information associated with <i>auth</i>. Destruction usually involves deallocation of private data structures. The use of <i>auth</i> is undefined after calling <code>auth_destroy()</code>.</pre> <p><code>authnone_create()</code> Create and return an RPC authentication handle that passes nonusable authentication information with each remote procedure call. This is the default authentication used by RPC.</p> <p><code>authsys_create()</code> Create and return an RPC authentication handle that contains AUTH_SYS authentication information. The parameter <i>host</i> is the name of the machine on which the information was created; <i>uid</i> is the user's user ID; <i>gid</i> is the user's current group ID; <i>len</i> and <i>aup_gids</i> refer to a counted array of groups to which the user belongs.</p> <p><code>authsys_create_default</code> Call <code>authsys_create()</code> with the appropriate parameters.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

rpc\_clnt\_auth(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `rpc(3NSL)`, `rpc_clnt_calls(3NSL)`, `rpc_clnt_create(3NSL)`,  
`secure_rpc(3NSL)`, `attributes(5)`

<b>NAME</b>	rpc_clnt_calls, clnt_call, clnt_send, clnt_freeres, clnt_geterr, clnt_perrno, clnt_perror, clnt_sperrno, clnt_sperror, rpc_broadcast, rpc_broadcast_exp, rpc_call – library routines for client side calls
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpc.h&gt;  enum clnt_stat <b>clnt_call</b>(CLIENT *clnt, const rpcproc_t procnum,     const xdrproc_t inproc, const caddr_t in, const xdrproc_t     outproc, caddr_t out, const struct timeval tout);  enum clnt_stat <b>clnt_send</b> (CLIENT *clnt, const u_long procnum, const     xdrproc_t proc, const caddr_t in);  bool_t <b>clnt_freeres</b>(CLIENT *clnt, const xdrproc_t outproc, caddr_t     out, );  void <b>clnt_geterr</b>(const CLIENT *clnt, struct rpc_err *errp);  void <b>clnt_perrno</b>(const enum clnt_stat stat);  void <b>clnt_perror</b>(const CLIENT *clnt, const char *s);  char *<b>clnt_sperrno</b>(const enum clnt_stat stat);  char *<b>clnt_sperror</b>(const CLIENT *clnt, const char *s);  enum clnt_stat <b>rpc_broadcast</b>(const rpcprog_t prognum, const     rpcvers_t versnum, const rpcproc_t procnum, const     xdrproc_t inproc, const caddr_t in, const xdrproc_t outproc,     caddr_t out, const resultproc_t eachresult, const char *nettype);  enum clnt_stat <b>rpc_broadcast_exp</b>(const rpcprog_t prognum, const     rpcvers_t versnum, const rpcproc_t procnum, const     xdrproc_t xargs, caddr_t argsp, const xdrproc_t txresults, caddr_t     resultsp, const resultproc_t eachresult, const int inittime, const int     waittime, const char *nettype);  enum clnt_stat <b>rpc_call</b>(const char *host, const rpcprog_t prognum,     const rpcvers_t versnum, const rpcproc_t procnum, const     xdrproc_t inproc, const char *in, const xdrproc_t outproc, char     *out, const char *nettype);</pre>
<b>DESCRIPTION</b>	<p>RPC library routines allow C language programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service and then sends back a reply.</p> <p>The <code>clnt_call()</code>, <code>rpc_call()</code>, and <code>rpc_broadcast()</code> routines handle the client side of the procedure call. The remaining routines deal with error handling.</p> <p>Some of the routines take a <code>CLIENT</code> handle as one of the parameters. A <code>CLIENT</code> handle can be created by an RPC creation routine such as <code>clnt_create()</code>. See <code>rpc_clnt_create(3NSL)</code>.</p>

## rpc\_clnt\_calls(3NSL)

These routines are safe for use in multithreaded applications. CLIENT handles can be shared between threads; however, in this implementation requests by different threads are serialized. In other words, the first request will receive its results before the second request is sent.

**Routines** See `rpc(3NSL)` for the definition of the CLIENT data structure.

### `clnt_call()`

A function macro that calls the remote procedure *procnum* associated with the client handle, *clnt*, which is obtained with an RPC client creation routine such as `clnt_create()`. See `rpc_clnt_create(3NSL)`. The parameter *inproc* is the XDR function used to encode the procedure's parameters, and *outproc* is the XDR function used to decode the procedure's results. *in* is the address of the procedure's argument(s), and *out* is the address of where to place the result(s). *tout* is the time allowed for results to be returned, which is overridden by a time-out set explicitly through `clnt_control()`. See `rpc_clnt_create(3NSL)`.

If the remote call succeeds, the status returned is `RPC_SUCCESS`. Otherwise, an appropriate status is returned.

### `clnt_send()`

Use the `clnt_send()` function to call a remote asynchronous function.

The `clnt_send()` function calls the remote function `procnum()` associated with the client handle, *clnt*, which is obtained with an RPC client creation routine such as `clnt_create()`. See `rpc_clnt_create(3NSL)`. The parameter *proc* is the XDR function used to encode the procedure's parameters. The parameter *in* is the address of the procedure's argument(s).

By default, the blocking I/O mode is used. See the `clnt_control(3NSL)` man page for more information on I/O modes.

The `clnt_send()` function does not check if the program version number supplied to `clnt_create()` is registered with the `rpcbind` service. Use `clnt_create_vers()` instead of `clnt_create()` to check on incorrect version number registration. `clnt_create_vers()` will return a valid handle to the client only if a version within the range supplied to `clnt_create_vers()` is supported by the server.

`RPC_SUCCESS` is returned when a request is successfully delivered to the transport layer. This does not mean that the request was received. If an error is returned, use the `clnt_getterr()` routine to find the failure status or the `clnt_perrno()` routine to translate the failure status into error messages.

### `clnt_freeres()`

A function macro that frees any data allocated by the RPC/XDR system when it decoded the results of an RPC call. The parameter *out* is the address of the results, and *outproc* is the XDR routine describing the results. This routine returns 1 if the results were successfully freed; otherwise it returns 0.



`clnt_geterr()`

A function macro that copies the error structure out of the client handle to the structure at address *errp*.

`clnt_perrno()`

Prints a message to standard error corresponding to the condition indicated by *stat*. A newline is appended. It is normally used after a procedure call fails for a routine for which a client handle is not needed, for instance `rpc_call()`.

`clnt_perror()`

Prints a message to the standard error indicating why an RPC call failed; *clnt* is the handle used to do the call. The message is prepended with string *s* and a colon. A newline is appended. This routine is normally used after a remote procedure call fails for a routine that requires a client handle, for instance `clnt_call()`.

`clnt_sperrno()`

Takes the same arguments as `clnt_perrno()`, but instead of sending a message to the standard error indicating why an RPC call failed, returns a pointer to a string that contains the message.

`clnt_sperrno()` is normally used instead of `clnt_perrno()` when the program does not have a standard error, as a program running as a server quite likely does not. `clnt_sperrno()` is also used if the programmer does not want the message to be output with `printf()`, or if a message format different than that supported by `clnt_perrno()` is to be used. See `printf(3C)`. Unlike `clnt_sperror()` and `clnt_sprecreateerror()`, `clnt_sperrno()` does not return a pointer to static data. Therefore, the result is not overwritten on each call. See `rpc_clnt_create(3NSL)`.

`clnt_sperror()`

Similar to `clnt_perror()`, except that like `clnt_sperrno()`, it returns a string instead of printing to standard error. However, `clnt_sperror()` does not append a newline at the end of the message.

`clnt_sperror()` returns a pointer to a buffer that is overwritten on each call. In multithreaded applications, this buffer is implemented as thread-specific data.

`rpc_broadcast()`

Similar to `rpc_call()`, except that the call message is broadcast to all the connectionless transports specified by *nettype*. If *nettype* is `NULL`, it defaults to `netpath`. Each time it receives a response, this routine calls `eachresult()`, whose form is:

```
bool_t eachresult(caddr_t out, const struct netbuf *addr,
const struct netconfig *netconf);
```

where *out* is the same as *out* passed to `rpc_broadcast()`, except that the remote procedure's output is decoded there. *addr* points to the address of the machine that sent the results, and *netconf* is the `netconfig` structure of the transport on which the remote server responded. If `eachresult()` returns 0, `rpc_broadcast()` waits for more replies; otherwise, it returns with appropriate status.

## rpc\_clnt\_calls(3NSL)

The broadcast file descriptors are limited in size to the maximum transfer size of that transport. For Ethernet, this value is 1500 bytes. `rpc_broadcast()` uses `AUTH_SYS` credentials by default. See `rpc_clnt_auth(3NSL)`.

### `rpc_broadcast_exp()`

Similar to `rpc_broadcast()`, except that the initial timeout, *inittime* and the maximum timeout, *waittime*, are specified in milliseconds.

*inittime* is the initial time that `rpc_broadcast_exp()` waits before resending the request. After the first resend, the retransmission interval increases exponentially until it exceeds *waittime*.

### `rpc_call()`

Calls the remote procedure associated with *prognum*, *versnum*, and *procnum* on the machine, *host*. The parameter *inproc* is used to encode the procedure's parameters, and *outproc* is used to decode the procedure's results. *in* is the address of the procedure's argument(s), and *out* is the address of where to place the result(s). *nettype* can be any of the values listed on `rpc(3NSL)`. This routine returns `RPC_SUCCESS` if it succeeds, or it returns an appropriate status. Use the `clnt_perrno()` routine to translate failure status into error messages.

The `rpc_call()` function uses the first available transport belonging to the class *nettype* on which it can create a connection. You do not have control of timeouts or authentication using this routine.

## ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	All
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Evolving
MT-Level	MT-Safe

## SEE ALSO

`printf(3C)`, `rpc(3NSL)`, `rpc_clnt_auth(3NSL)`, `rpc_clnt_create(3NSL)`, `attributes(5)`

**NAME** | rpc\_clnt\_create, clnt\_control, clnt\_create, clnt\_create\_timed, clnt\_create\_vers, clnt\_create\_vers\_timed, clnt\_destroy, clnt\_dg\_create, clnt\_pcreateerror, clnt\_raw\_create, clnt\_spcreateerror, clnt\_tli\_create, clnt\_tp\_create, clnt\_tp\_create\_timed, clnt\_vc\_create, rpc\_createerr, clnt\_door\_create – library routines for dealing with creation and manipulation of CLIENT handles

**SYNOPSIS** | #include <rpc/rpc.h>

```

bool_t clnt_control(CLIENT *clnt, const uint_t req, char *info);

CLIENT *clnt_create(const char *host, const rpcprog_t prognum,
    const rpcvers_t versnum, const char *nettype);

CLIENT *clnt_create_timed(const char *host, const rpcprog_t
    prognum, const rpcvers_t versnum, const nettype, const struct
    timeval *timeout);

CLIENT *clnt_create_vers(const char *host, const rpcprog_t prognum,
    rpcvers_t *vers_outp, const rpcvers_t vers_low, const rpcvers_t
    vers_high, char *nettype);

CLIENT *clnt_create_vers_timed(const char *host, const rpcprog_t
    prognum, rpcvers_t *vers_outp, const rpcvers_t vers_low, const
    rpcvers_t vers_high, char *nettype, const struct timeval *timeout);

void clnt_destroy(CLIENT *clnt);

CLIENT *clnt_dg_create(const int fildes, const struct netbuf
    *svcaddr, const rpcprog_t prognum, const rpcvers_t versnum, const
    uint_t sendsz, const uint_t recsz);

void clnt_pcreateerror(const char *s);

CLIENT *clnt_raw_create(const rpcprog_t prognum, const rpcvers_t
    versnum);

char *clnt_spcreateerror(const char *s);

CLIENT *clnt_tli_create(const int fildes, const struct netconfig
    *netconf, const struct netbuf *svcaddr, const rpcprog_t prognum,
    const rpcvers_t versnum, const uint_t sendsz, const uint_t
    recsz);

CLIENT *clnt_tp_create(const char *host, const rpcprog_t prognum,
    const rpcvers_t versnum, const struct netconfig *netconf);

CLIENT *clnt_tp_create_timed(const char *host, const rpcprog_t
    prognum, const rpcvers_t versnum, const struct netconfig
    *netconf, const struct timeval *timeout);

CLIENT *clnt_vc_create(const int fildes, const struct netbuf
    *svcaddr, const rpcprog_t prognum, const rpcvers_t versnum, const
    uint_t sendsz, const uint_t recsz);

struct rpc_createerr rpc_createerr

```

## rpc\_clnt\_create(3NSL)

	<pre>CLIENT *clnt_door_create(const rpcprog_t prognum, const rpcvers_t     versnum, const uint_t sendsz);</pre>
<b>DESCRIPTION</b>	<p>RPC library routines allow C language programs to make procedure calls on other machines across the network. First a CLIENT handle is created and then the client calls a procedure to send a request to the server. On receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends a reply.</p> <p>These routines are MT-Safe. In the case of multithreaded applications, the <code>-mt</code> option must be specified on the command line at compilation time. When the <code>-mt</code> option is specified, <code>rpc_createerr()</code> becomes a macro that enables each thread to have its own <code>rpc_createerr()</code>. See <code>threads(3THR)</code>.</p>
<b>Routines</b>	<p>See <code>rpc(3NSL)</code> for the definition of the CLIENT data structure.</p> <pre>clnt_control()</pre> <p>A function macro to change or retrieve various information about a client object. <i>req</i> indicates the type of operation, and <i>info</i> is a pointer to the information. For both connectionless and connection-oriented transports, the supported values of <i>req</i> and their argument types and what they do are:</p> <pre>CLSET_TIMEOUT struct timeval * set total timeout CLGET_TIMEOUT  struct timeval *  get total timeout</pre> <p>If the timeout is set using <code>clnt_control()</code>, the timeout argument passed by <code>clnt_call()</code> is ignored in all subsequent calls. If the timeout value is set to 0, <code>clnt_control()</code> immediately returns <code>RPC_TIMEDOUT</code>. Set the timeout parameter to 0 for batching calls.</p> <pre>CLGET_SERVER_ADDR struct netbuf * get server's address CLGET_SVC_ADDR    struct netbuf *  get server's address CLGET_FD          int *           get associated file descriptor CLSET_FD_CLOSE    void           close the file descriptor when                                 destroying the client handle                                 (see clnt_destroy()) CLSET_FD_NCLOSE   void           do not close the file                                 descriptor when destroying the client handle CLGET_VERS        rpcvers_t      get the RPC program's version                                 number associated with the                                 client handle CLSET_VERS        rpcvers_t      set the RPC program's version                                 number associated with the                                 client handle. This assumes                                 that the RPC server for this                                 new version is still listening                                 at the address of the previous                                 version. CLGET_XID         uint32_t       get the XID of the previous                                 remote procedure call CLSET_XID         uint32_t       set the XID of the next                                 remote procedure call CLGET_PROG        rpcprog_t      get program number CLSET_PROG        rpcprog_t      set program number</pre> <p>The following operations are valid for connection-oriented transports only:</p>

## rpc\_clnt\_create(3NSL)

`CLSET_IO_MODE` `rpciomode_t*` set the IO mode used to send one-way requests. The argument for this operation can be either:

- `RPC_CL_BLOCKING` all sending operations block until the underlying transport protocol has accepted requests. If you specify this argument you cannot use flush and getting and setting buffer size is meaningless.
- `RPC_CL_NONBLOCKING` sending operations do not block and return as soon as requests enter the buffer. You can now use non-blocking I/O. The requests in the buffer are pending. The requests are sent to the server as soon as a two-way request is sent or a flush is done. You are responsible for flushing the buffer. When you choose `RPC_CL_NONBLOCKING` argument you have a choice of flush modes as specified by `CLSET_FLUSH_MODE`.

`CLGET_IO_MODE` `rpciomode_t*` get the current IO mode

`CLSET_FLUSH_MODE` `rpcflushmode_t*` set the flush mode. The flush mode can only be used in non-blocking I/O mode. The argument can be either of the following:

- `RPC_CL_BESTEFFORT_FLUSH`: All flushes send requests in the buffer until the transport end-point blocks. If the transport connection is congested, the call returns directly.
- `RPC_CL_BLOCKING_FLUSH`: Flush blocks until the underlying transport protocol accepts all pending requests into the queue.

`CLGET_FLUSH_MODE` `rpcflushmode_t*` get the current flush mode.

`CLFLUSH` `rpcflushmode_t` flush the pending requests. This command can only be used in non-blocking I/O mode. The flush policy depends on which of the following parameters is specified:

- `RPC_CL_DEFAULT_FLUSH`, or `NULL`: The flush is done according to the current flush mode policy (see `CLSET_FLUSH_MODE` option).
- `RPC_CL_BESTEFFORT_FLUSH`: The flush tries to send pending requests without blocking; the call returns directly. If the transport connection is congested, this call could return without the request being sent.
- `RPC_CL_BLOCKING_FLUSH`: The flush sends all pending requests. This call will block until all the requests have been accepted by the transport layer.

`CLSET_CONNMAXREC_SIZE` `int*` set the buffer size. It is not possible to dynamically resize the buffer if it contains data. The default size of the buffer is 16 kilobytes.

`CLGET_CONNMAXREC_SIZE` `int*` get the current size of the buffer

`CLGET_CURRENT_REC_SIZE` `int*` get the size of the pending requests stored in the buffer. Use of this command is only recommended when you are in non-blocking I/O mode. The current size of the buffer is always zero when the handle is in blocking mode as the buffer is not used in this mode.

The following operations are valid for connectionless transports only:

## rpc\_clnt\_create(3NSL)

```
CLSET_RETRY_TIMEOUT struct timeval *   set the retry timeout
CLGET_RETRY_TIMEOUT struct timeval *   get the retry timeout
```

The retry timeout is the time that RPC waits for the server to reply before retransmitting the request.

`clnt_control()` returns TRUE on success and FALSE on failure.

### `clnt_create()`

Generic client creation routine for program *prognum* and version *versnum*. *host* identifies the name of the remote host where the server is located. *nettype* indicates the class of transport protocol to use. The transports are tried in left to right order in NETPATH variable or in top to bottom order in the netconfig database.

`clnt_create()` tries all the transports of the *nettype* class available from the NETPATH environment variable and the netconfig database, and chooses the first successful one. A default timeout is set and can be modified using `clnt_control()`. This routine returns NULL if it fails. The `clnt_pcreateerror()` routine can be used to print the reason for failure.

Note that `clnt_create()` returns a valid client handle even if the particular version number supplied to `clnt_create()` is not registered with the `rpcbind` service. This mismatch will be discovered by a `clnt_call` later (see `rpc_clnt_calls(3NSL)`).

### `clnt_create_timed()`

Generic client creation routine which is similar to `clnt_create()` but which also has the additional parameter *timeout* that specifies the maximum amount of time allowed for each transport class tried. In all other respects, the `clnt_create_timed()` call behaves exactly like the `clnt_create()` call.

### `clnt_create_vers()`

Generic client creation routine which is similar to `clnt_create()` but which also checks for the version availability. *host* identifies the name of the remote host where the server is located. *nettype* indicates the class transport protocols to be used. If the routine is successful it returns a client handle created for the highest version between *vers\_low* and *vers\_high* that is supported by the server. *vers\_outp* is set to this value. That is, after a successful return `vers_low <= *vers_outp <= vers_high`. If no version between *vers\_low* and *vers\_high* is supported by the server then the routine fails and returns NULL. A default timeout is set and can be modified using `clnt_control()`. This routine returns NULL if it fails. The `clnt_pcreateerror()` routine can be used to print the reason for failure.

Note: `clnt_create()` returns a valid client handle even if the particular version number supplied to `clnt_create()` is not registered with the `rpcbind` service. This mismatch will be discovered by a `clnt_call` later (see `rpc_clnt_calls(3NSL)`). However, `clnt_create_vers()` does this for you and returns a valid handle only if a version within the range supplied is supported by the server.

`clnt_create_vers_timed()`

Generic client creation routine similar to `clnt_create_vers()` but with the additional parameter *timeout*, which specifies the maximum amount of time allowed for each transport class tried. In all other respects, the `clnt_create_vers_timed()` call behaves exactly like the `clnt_create_vers()` call.

`clnt_destroy()`

A function macro that destroys the client's RPC handle. Destruction usually involves deallocation of private data structures, including *clnt* itself. Use of *clnt* is undefined after calling `clnt_destroy()`. If the RPC library opened the associated file descriptor, or `CLSET_FD_CLOSE` was set using `clnt_control()`, the file descriptor will be closed.

The caller should call `auth_destroy(clnt->cl_auth)` (before calling `clnt_destroy()`) to destroy the associated AUTH structure (see `rpc_clnt_auth(3NSL)`).

`clnt_dg_create()`

This routine creates an RPC client for the remote program *prognum* and version *versnum*; the client uses a connectionless transport. The remote program is located at address *svcaddr*. The parameter *fdes* is an open and bound file descriptor. This routine will resend the call message in intervals of 15 seconds until a response is received or until the call times out. The total time for the call to time out is specified by `clnt_call()` (see `clnt_call()` in `rpc_clnt_calls(3NSL)`). The retry time out and the total time out periods can be changed using `clnt_control()`. The user may set the size of the send and receive buffers with the parameters *sendsz* and *rcvsz*; values of 0 choose suitable defaults. This routine returns NULL if it fails.

`clnt_pcreateerror()`

Print a message to standard error indicating why a client RPC handle could not be created. The message is prepended with the string *s* and a colon, and appended with a newline.

`clnt_raw_create()`

This routine creates an RPC client handle for the remote program *prognum* and version *versnum*. The transport used to pass messages to the service is a buffer within the process's address space, so the corresponding RPC server should live in the same address space; (see `svc_raw_create()` in `rpc_svc_create(3NSL)`). This allows simulation of RPC and measurement of RPC overheads, such as round trip times, without any kernel or networking interference. This routine returns NULL if it fails. `clnt_raw_create()` should be called after `svc_raw_create()`.

`clnt_spcreateerror()`

Like `clnt_pcreateerror()`, except that it returns a string instead of printing to the standard error. A newline is not appended to the message in this case.

Warning: returns a pointer to a buffer that is overwritten on each call. In multithread applications, this buffer is implemented as thread-specific data.

## rpc\_clnt\_create(3NSL)

### clnt\_tli\_create()

This routine creates an RPC client handle for the remote program *prognum* and version *versnum*. The remote program is located at address *svcaddr*. If *svcaddr* is NULL and it is connection-oriented, it is assumed that the file descriptor is connected. For connectionless transports, if *svcaddr* is NULL, RPC\_UNKNOWNADDR error is set. *fildev* is a file descriptor which may be open, bound and connected. If it is RPC\_ANYFD, it opens a file descriptor on the transport specified by *netconf*. If *fildev* is RPC\_ANYFD and *netconf* is NULL, a RPC\_UNKNOWNPROTO error is set. If *fildev* is unbound, then it will attempt to bind the descriptor. The user may specify the size of the buffers with the parameters *sendsz* and *recvsz*; values of 0 choose suitable defaults. Depending upon the type of the transport (connection-oriented or connectionless), *clnt\_tli\_create()* calls appropriate client creation routines. This routine returns NULL if it fails. The *clnt\_pcreateerror()* routine can be used to print the reason for failure. The remote *rpcbind* service (see *rpcbind(1M)*) is not consulted for the address of the remote service.

### clnt\_tp\_create()

Like *clnt\_create()* except *clnt\_tp\_create()* tries only one transport specified through *netconf*.

*clnt\_tp\_create()* creates a client handle for the program *prognum*, the version *versnum*, and for the transport specified by *netconf*. Default options are set, which can be changed using *clnt\_control()* calls. The remote *rpcbind* service on the host *host* is consulted for the address of the remote service. This routine returns NULL if it fails. The *clnt\_pcreateerror()* routine can be used to print the reason for failure.

### clnt\_tp\_create\_timed()

Like *clnt\_tp\_create()* except *clnt\_tp\_create\_timed()* has the extra parameter *timeout* which specifies the maximum time allowed for the creation attempt to succeed. In all other respects, the *clnt\_tp\_create\_timed()* call behaves exactly like the *clnt\_tp\_create()* call.

### clnt\_vc\_create()

This routine creates an RPC client for the remote program *prognum* and version *versnum*; the client uses a connection-oriented transport. The remote program is located at address *svcaddr*. The parameter *fildev* is an open and bound file descriptor. The user may specify the size of the send and receive buffers with the parameters *sendsz* and *recvsz*; values of 0 choose suitable defaults. This routine returns NULL if it fails.

The address *svcaddr* should not be NULL and should point to the actual address of the remote program. *clnt\_vc\_create()* does not consult the remote *rpcbind* service for this information.

### rpc\_createerr()

A global variable whose value is set by any RPC client handle creation routine that fails. It is used by the routine *clnt\_pcreateerror()* to print the reason for the failure.



rpc\_clnt\_create(3NSL)

In multithreaded applications, `rpc_createerr` becomes a macro which enables each thread to have its own `rpc_createerr`.

`clnt_door_create()`

This routine creates an RPC client handle over doors for the given program *prognum* and version *versnum*. Doors is a transport mechanism that facilitates fast data transfer between processes on the same machine. The user may set the size of the send buffer with the parameter *sendsz*. If *sendsz* is 0, the corresponding default buffer size is 16 Kbyte. The `clnt_door_create()` routine returns `NULL` if it fails and sets a value for `rpc_createerr`.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	All
Availability	SUNWcsl (32-bit) SUNWcslx (64-bit)
Interface Stability	Evolving
MT-Level	MT-Safe

**SEE ALSO** `rpcbind(1M)`, `rpc(3NSL)`, `rpc_clnt_auth(3NSL)`, `rpc_clnt_calls(3NSL)`, `rpc_svc_create(3NSL)`, `svc_raw_create(3NSL)`, `threads(3THR)`, `attributes(5)`

rpc\_control(3NSL)

**NAME** | rpc\_control – library routine for manipulating global RPC attributes for client and server applications

**SYNOPSIS** | `bool_t rpc_control(int op, void *info);`

**DESCRIPTION** | This RPC library routine allows applications to set and modify global RPC attributes that apply to clients as well as servers. At present, it supports only server side operations. This function allows applications to set and modify global attributes that apply to client as well as server functions. *op* indicates the type of operation, and *info* is a pointer to the operation specific information. The supported values of *op* and their argument types, and what they do are:

RPC_SVC_MTMODE_SET	int *	set multithread mode
RPC_SVC_MTMODE_GET	int *	get multithread mode
RPC_SVC_THRMAX_SET	int *	set maximum number of threads
RPC_SVC_THRMAX_GET	int *	get maximum number of threads
RPC_SVC_THRTOTAL_GET	int *	get number of active threads
RPC_SVC_THRCREATES_GET	int *	get number of threads created
RPC_SVC_THRERRORS_GET	int *	get number of thread create errors
RPC_SVC_USE_POLLFD	int *	set number of file descriptors to unlimited
RPC_SVC_CONNMAXREC_SET	int *	set non-blocking max rec size
RPC_SVC_CONNMAXREC_GET	int *	get non-blocking max rec size

There are three multithread (MT) modes. These are:

RPC_SVC_MT_NONE	Single threaded mode	(default)
RPC_SVC_MT_AUTO	Automatic MT mode	
RPC_SVC_MT_USER	User MT mode	

Unless the application sets the Automatic or User MT modes, it will stay in the default (single threaded) mode. See the *Network Interfaces Programmer's Guide* for the meanings of these modes and programming examples. Once a mode is set, it cannot be changed.

By default, the maximum number of threads that the server will create at any time is 16. This allows the service developer to put a bound on thread resources consumed by a server. If a server needs to process more than 16 client requests concurrently, the maximum number of threads must be set to the desired number. This parameter may be set at any time by the server.

Set and get operations will succeed even in modes where the operations don't apply. For example, you can set the maximum number of threads in any mode, even though it makes sense only for the Automatic MT mode. All of the get operations except `RPC_SVC_MTMODE_GET` apply only to the Automatic MT mode, so values returned in other modes may be undefined.

By default, RPC servers are limited to a maximum of 1024 file descriptors or connections due to limitations in the historical interfaces `svc_fdset(3NSL)` and `svc_getreqset(3NSL)`. Applications written to use the preferred interfaces of `svc_pollfd(3NSL)` and `svc_getreq_poll(3NSL)` can use an unlimited number of file descriptors. Setting *info* to point to a non-zero integer and *op* to `RPC_SVC_USE_POLLFD` removes the limitation.

rpc\_control(3NSL)

Connection oriented RPC transports read RPC requests in blocking mode by default. Thus, they may be adversely affected by network delays and broken clients. `RPC_SVC_CONNMAXREC_SET` enables non-blocking mode and establishes the maximum record size (in bytes) for RPC requests; RPC responses are not affected. Buffer space is allocated as needed up to the specified maximum, starting at the maximum or `RPC_MAXDATASIZE`, whichever is smaller.

The value established by `RPC_SVC_CONNMAXREC_SET` is used when a connection is created, and it remains in effect for that connection until it is closed. To change the value for existing connections on a per-connection basis, see `svc_control(3NSL)`.

`RPC_SVC_CONNMAXREC_GET` retrieves the current maximum record size. A zero value means that no maximum is in effect, and that the connections are in blocking mode.

*info* is a pointer to an argument of type `int`. Non-connection RPC transports ignore `RPC_SVC_CONNMAXREC_SET` and `RPC_SVC_CONNMAXREC_GET`.

**RETURN VALUES** This routine returns `TRUE` if the operation was successful and returns `FALSE` otherwise.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `rpcbind(1M)`, `rpc(3NSL)`, `rpc_svc_calls(3NSL)`, `attributes(5)`

*Network Interfaces Programmer's Guide*

## rpc\_gss\_getcred(3NSL)

<b>NAME</b>	rpc_gss_getcred – get credentials of client
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpcsec_gss.h&gt;  bool_t <b>rpc_gss_getcred</b>(struct svc_req *req, rpc_gss_rawcred_t     **rcred, rpc_gss_ucred **ucred, void **cookie);</pre>
<b>DESCRIPTION</b>	<p>rpc_gss_getcred() is used by a server to fetch the credentials of a client. These credentials may either be network credentials (in the form of a rpc_gss_rawcred_t structure) or UNIX credentials.</p> <p>For more information on RPCSEC_GSS data types, see the rpcsec_gss(3NSL) man page.</p>
<b>PARAMETERS</b>	<p>Essentially, rpc_gss_getcred() passes a pointer to a request (svc_req) as well as pointers to two credential structures and a user-defined cookie; if rpc_gss_getcred() is successful, at least one credential structure is "filled out" with values, as is, optionally, the cookie.</p> <p><i>req</i>                      Pointer to the received service request. svc_req is an RPC structure containing information on the context of an RPC invocation, such as program, version, and transport information.</p> <p><i>rcred</i>                     A pointer to an rpc_gss_rawcred_t structure pointer. This structure contains the version number of the RPCSEC_GSS protocol being used; the security mechanism and QOPs for this session (as strings); principal names for the client (as a rpc_gss_principal_t structure) and server (as a string); and the security service (integrity, privacy, etc., as an enum). If an application is not interested in these values, it may pass NULL for this parameter.</p> <p><i>ucred</i>                    The caller's UNIX credentials, in the form of a pointer to a pointer to a rpc_gss_ucred_t structure, which includes the client's uid and gids. If an application is not interested in these values, it may pass NULL for this parameter.</p> <p><i>cookie</i>                   A four-byte quantity that an application may use in any manner it wants to; RPC does not interpret it. (For example, a cookie may be a pointer or index to a structure that represents a context initiator.) See also rpc_gss_set_callback(3NSL).</p>
<b>RETURN VALUES</b>	rpc_gss_getcred() returns TRUE if it is successful; otherwise, use rpc_gss_get_error() to get the error associated with the failure.
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

rpc\_gss\_getcred(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWrsg (32-bit)
	SUNWrsgx (64-bit)

**SEE ALSO** `rpc(3NSL)`, `rpc_gss_set_callback(3NSL)`, `rpc_gss_set_svc_name(3NSL)`, `rpcsec_gss(3NSL)`, `attributes(5)`

*ONC+ Developer's Guide*

Linn, J. *RFC 2078, Generic Security Service Application Program Interface, Version 2*. Network Working Group. January 1997.

## rpc\_gss\_get\_error(3NSL)

<b>NAME</b>	rpc_gss_get_error – get error codes on failure								
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpcsec_gss.h&gt;  bool_t <b>rpc_gss_get_error</b>(rpc_gss_error_t*error);</pre>								
<b>DESCRIPTION</b>	<p>rpc_gss_get_error() fetches an error code when an RPCSEC_GSS routine fails.</p> <p>rpc_gss_get_error() uses a rpc_gss_error_t structure of the following form:</p> <pre>typedef struct {     int    rpc_gss_error;          <i>RPCSEC_GSS error</i>     int    system_error;          <i>system error</i> } rpc_gss_error_t;</pre> <p>Currently the only error codes defined for this function are</p> <pre>#define RPC_GSS_ER_SUCCESS      0    /* no error */ #define RPC_GSS_ER_SYSTEMERROR  1    /* system error */</pre>								
<b>PARAMETERS</b>	<p>Information on RPCSEC_GSS data types for parameters may be found on the rpcsec_gss(3NSL) man page.</p> <p>error                   A rpc_gss_error_t structure. If the rpc_gss_error field is equal to RPC_GSS_ER_SYSTEMERROR, the system_error field will be set to the value of errno.</p>								
<b>RETURN VALUES</b>	Unless there is a failure indication from an invoked RPCSEC_GSS function, rpc_gss_get_error() does not set error to a meaningful value.								
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:								
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr><tr><td>Availability</td><td>SUNWrsg (32-bit)</td></tr><tr><td></td><td>SUNWrsgx (64-bit)</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe	Availability	SUNWrsg (32-bit)		SUNWrsgx (64-bit)
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
MT-Level	MT-Safe								
Availability	SUNWrsg (32-bit)								
	SUNWrsgx (64-bit)								
<b>SEE ALSO</b>	<p>perror(3C), rpc(3NSL), rpcsec_gss(3NSL), attributes(5)</p> <p><i>ONC+ Developer's Guide</i></p> <p>Linn, J. <i>RFC 2078, Generic Security Service Application Program Interface, Version 2.</i> Network Working Group. January 1997.</p>								
<b>NOTES</b>	Only system errors are currently returned.								

<b>NAME</b>	rpc_gss_get_mechanisms, rpc_gss_get_mech_info, rpc_gss_get_versions, rpc_gss_is_installed – get information on mechanisms and RPC version
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpcsec_gss.h&gt;  char **rpc_gss_get_mechanisms();  char **rpc_gss_get_mech_info(char *mech, rpc_gss_service_t     *service);  bool_t rpc_gss_get_versions(u_int *vers_hi, u_int *vers_lo);  bool_t rpc_gss_is_installed(char *mech);</pre>
<b>DESCRIPTION</b>	<p>These "convenience functions" return information on available security mechanisms and versions of RPCSEC_GSS.</p> <p>rpc_gss_get_mechanisms() Returns a list of supported security mechanisms as a null-terminated list of character strings.</p> <p>rpc_gss_get_mech_info() Takes two arguments: an ASCII string representing a mechanism type, for example, kerberosv5, and a pointer to a rpc_gss_service_t enum. rpc_gss_get_mech_info() will return NULL upon error or if no /etc/gss/qop file is present. Otherwise, it returns a null-terminated list of character strings of supported Quality of Protections (QOPs) for this mechanism. NULL or empty list implies only that the default QOP is available and can be specified to routines that need to take a QOP string parameter as NULL or as an empty string.</p> <p>rpc_gss_get_versions() Returns the highest and lowest versions of RPCSEC_GSS supported.</p> <p>rpc_gss_is_installed() Takes an ASCII string representing a mechanism, and returns TRUE if the mechanism is installed.</p>
<b>PARAMETERS</b>	<p>Information on RPCSEC_GSS data types for parameters may be found on the rpcsec_gss(3NSL) man page.</p> <p><i>mech</i> An ASCII string representing the security mechanism in use. Valid strings may also be found in the /etc/gss/mech file.</p> <p><i>service</i> A pointer to a rpc_gss_service_t enum, representing the current security service (privacy, integrity, or none).</p>

## rpc\_gss\_get\_mechanisms(3NSL)

	<i>vers_hi</i>									
	<i>vers_lo</i>	The highest and lowest versions of RPCSEC_GSS supported.								
<b>FILES</b>	/etc/gss/mech	File containing valid security mechanisms								
	/etc/gss/qop	File containing valid QOP values								
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:									
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr><tr><td>Availability</td><td>SUNWrsg (32-bit)</td></tr><tr><td></td><td>SUNWrsgx (64-bit)</td></tr></tbody></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe	Availability	SUNWrsg (32-bit)		SUNWrsgx (64-bit)
ATTRIBUTE TYPE	ATTRIBUTE VALUE									
MT-Level	MT-Safe									
Availability	SUNWrsg (32-bit)									
	SUNWrsgx (64-bit)									
<b>SEE ALSO</b>	<code>rpc(3NSL)</code> , <code>rpcsec_gss(3NSL)</code> , <code>mech(4)</code> , <code>qop(4)</code> , <code>attributes(5)</code> <i>ONC+ Developer's Guide</i> Linn, J. <i>RFC 2743, Generic Security Service Application Program Interface Version 2, Update 1</i> . Network Working Group. January 2000.									
<b>NOTES</b>	This function will change in a future release.									



rpc\_gss\_get\_principal\_name(3NSL)

<b>NAME</b>	rpc_gss_get_principal_name – Get principal names at server										
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpcsec_gss.h&gt;  bool_t <b>rpc_gss_get_principal_name</b>(rpc_gss_principal_ *principal,     char *mech, char *name, char *node, char *domain);</pre>										
<b>DESCRIPTION</b>	<p>Servers need to be able to operate on a client’s principal name. Such a name is stored by the server as a <code>rpc_gss_principal_t</code> structure, an opaque byte string which can be used either directly in access control lists or as database indices which can be used to look up a UNIX credential. A server may, for example, need to compare a principal name it has received with the principal name of a known entity, and to do that, it must be able to generate <code>rpc_gss_principal_t</code> structures from known entities.</p> <p><code>rpc_gss_get_principal_name()</code> takes as input a security mechanism, a pointer to a <code>rpc_gss_principal_t</code> structure, and several parameters which uniquely identify an entity on a network: a user or service name, a node name, and a domain name. From these parameters it constructs a unique, mechanism-dependent principal name of the <code>rpc_gss_principal_t</code> structure type.</p>										
<b>PARAMETERS</b>	<p>How many of the identifying parameters (<i>name</i>, <i>node</i>, and <i>domain</i>) are necessary to specify depends on the mechanism being used. For example, Kerberos V5 requires only a user name but can accept a node and domain name. An application can choose to set unneeded parameters to <code>NULL</code>.</p> <p>Information on <code>RPCSEC_GSS</code> data types for parameters may be found on the <code>rpcsec_gss(3NSL)</code> man page.</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><i>principal</i></td> <td>An opaque, mechanism-dependent structure representing the client’s principal name.</td> </tr> <tr> <td><i>mech</i></td> <td>An ASCII string representing the security mechanism in use. Valid strings may be found in the <code>/etc/gss/mech</code> file, or by using <code>rpc_gss_get_mechanisms()</code>.</td> </tr> <tr> <td><i>name</i></td> <td>A UNIX login name (for example, ‘<code>gwashington</code>’) or service name, such as ‘<code>nfs</code>’.</td> </tr> <tr> <td><i>node</i></td> <td>A node in a domain; typically, this would be a machine name (for example, ‘<code>valleyforge</code>’).</td> </tr> <tr> <td><i>domain</i></td> <td>A security domain; for example, a DNS, NIS, or NIS+ domain name (‘<code>eng.company.com</code>’).</td> </tr> </table>	<i>principal</i>	An opaque, mechanism-dependent structure representing the client’s principal name.	<i>mech</i>	An ASCII string representing the security mechanism in use. Valid strings may be found in the <code>/etc/gss/mech</code> file, or by using <code>rpc_gss_get_mechanisms()</code> .	<i>name</i>	A UNIX login name (for example, ‘ <code>gwashington</code> ’) or service name, such as ‘ <code>nfs</code> ’.	<i>node</i>	A node in a domain; typically, this would be a machine name (for example, ‘ <code>valleyforge</code> ’).	<i>domain</i>	A security domain; for example, a DNS, NIS, or NIS+ domain name (‘ <code>eng.company.com</code> ’).
<i>principal</i>	An opaque, mechanism-dependent structure representing the client’s principal name.										
<i>mech</i>	An ASCII string representing the security mechanism in use. Valid strings may be found in the <code>/etc/gss/mech</code> file, or by using <code>rpc_gss_get_mechanisms()</code> .										
<i>name</i>	A UNIX login name (for example, ‘ <code>gwashington</code> ’) or service name, such as ‘ <code>nfs</code> ’.										
<i>node</i>	A node in a domain; typically, this would be a machine name (for example, ‘ <code>valleyforge</code> ’).										
<i>domain</i>	A security domain; for example, a DNS, NIS, or NIS+ domain name (‘ <code>eng.company.com</code> ’).										
<b>RETURN VALUES</b>	<code>rpc_gss_get_principal_name()</code> returns <code>TRUE</code> if it is successful; otherwise, use <code>rpc_gss_get_error()</code> to get the error associated with the failure.										
<b>FILES</b>	<code>/etc/gss/mech</code> File containing valid security mechanisms										
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:										

rpc\_gss\_get\_principal\_name(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe
Availability	SUNWrsg (32-bit)
	SUNWrsgx (64-bit)

**SEE ALSO** free(3C), rpc(3NSL), rpc\_gss\_get\_mechanisms(3NSL),  
rpc\_gss\_set\_svc\_name(3NSL), rpcsec\_gss(3NSL), mech(4), attributes(5)

*ONC+ Developer's Guide*

Linn, J. *RFC 2078, Generic Security Service Application Program Interface, Version 2.*  
Network Working Group. January 1997.

**NOTES** Principal names may be freed up by a call to `free(3C)`. A principal name need only be freed in those instances where it was constructed by the application. (Values returned by other routines point to structures already existing in a context, and need not be freed.)

rpc\_gss\_max\_data\_length(3NSL)

**NAME** | `rpc_gss_max_data_length`, `rpc_gss_svc_max_data_length` – get maximum data length for transmission

**SYNOPSIS** | 

```
#include <rpc/rpcsec_gss.h>

int rpc_gss_max_data_length(AUTH *handle, int max_tp_unit_len);

int rpc_gss_svc_max_data_length(struct svc_req *req, int
    max_tp_unit_len);
```

**DESCRIPTION** | Performing a security transformation on a piece of data generally produces data with a different (usually greater) length. For some transports, such as UDP, there is a maximum length of data which can be sent out in one data unit. Applications need to know the maximum size a piece of data can be before it's transformed, so that the resulting data will still "fit" on the transport. These two functions return that maximum size.

| `rpc_gss_max_data_length()` is the client-side version;  
| `rpc_gss_svc_max_data_length()` is the server-side version.

**PARAMETERS** | *handle* | An RPC context handle of type AUTH, returned when a context is created (for example, by `rpc_gss_seccreate()`). Security service and QOP are bound to this handle, eliminating any need to specify them.

| *max\_tp\_unit\_len* | The maximum size of a piece of data allowed by the transport.

| *req* | A pointer to an RPC `svc_req` structure, containing information on the context (for example, program number and credentials).

**RETURN VALUES** | Both functions return the maximum size of untransformed data allowed, as an int.

**ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe
Availability	SUNWrsg (32-bit)
	SUNWrsgx (64-bit)

**SEE ALSO** | `rpc(3NSL)`, `rpcsec_gss(3NSL)`, `attributes(5)`

| *ONC+ Developer's Guide*

| Linn, J. *RFC 2078, Generic Security Service Application Program Interface, Version 2.* Network Working Group. January 1997.

rpc\_gss\_mech\_to\_oid(3NSL)

**NAME** | rpc\_gss\_mech\_to\_oid, rpc\_gss\_qop\_to\_num – map mechanism, QOP strings to non-string values

**SYNOPSIS** | #include <rpc/rpcsec\_gss.h>  
 | bool\_t **rpc\_gss\_mech\_to\_oid**(charc\*mech, rpc\_gss\_OIDc\*oid);  
 | bool\_t **rpc\_gss\_qop\_to\_num**(char \*qop, char \*mech, u\_int \*num);

**DESCRIPTION** | Because in-kernel RPC routines use non-string values for mechanism and Quality of Protection (QOP), these routines exist to map strings for these attributes to their non-string counterparts. (The non-string values for QOP and mechanism are also found in the /etc/gss/qop and /etc/gss/mech files, respectively.)  
 | rpc\_gss\_mech\_to\_oid() takes a string representing a mechanism, as well as a pointer to a rpc\_gss\_OID object identifier structure. It then gives this structure values corresponding to the indicated mechanism, so that the application can now use the OID directly with RPC routines. rpc\_gss\_qop\_to\_num() does much the same thing, taking strings for QOP and mechanism and returning a number.

**PARAMETERS** | Information on RPCSEC\_GSS data types for parameters may be found on the rpcsec\_gss(3NSL) man page.

*mech* | An ASCII string representing the security mechanism in use. Valid strings may be found in the /etc/gss/mech file.

*oid* | An object identifier of type rpc\_gss\_OID, whose elements are usable by kernel-level RPC routines.

*qop* | This is an ASCII string which sets the quality of protection (QOP) for the session. Appropriate values for this string may be found in the file /etc/gss/qop.

*num* | The non-string value for the QOP.

**RETURN VALUES** | Both functions return TRUE if they are successful, FALSE otherwise.

**FILES** | /etc/gss/mech | File containing valid security mechanisms  
 | /etc/gss/qop | File containing valid QOP values

**ATTRIBUTES** | See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe
Availability	SUNWrsg (32-bit)
	SUNWrsgx (64-bit)

**SEE ALSO** | rpc(3NSL), rpc\_gss\_get\_error(3NSL), rpc\_gss\_get\_mechanisms(3NSL), rpcsec\_gss(3NSL), mech(4), qop(4), attributes(5)

rpc\_gss\_mech\_to\_oid(3NSL)

*ONC+ Developer's Guide*

Linn, J. *RFC 2078, Generic Security Service Application Program Interface, Version 2.*  
Network Working Group. January 1997.

## rpc\_gss\_seccreate(3NSL)

<b>NAME</b>	rpc_gss_seccreate – create a security context using the RPCSEC_GSS protocol														
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpcsec_gss.h&gt;  AUTH *rpc_gss_seccreate(CLIENT *clnt, char *principal, char *mechanism,     rpc_gss_service_t service_type, char *qop, rpc_gss_options_req_t     *options_req, rpc_gss_options_ret_t *options_ret);</pre>														
<b>DESCRIPTION</b>	<p>rpc_gss_seccreate() is used by an application to create a security context using the RPCSEC_GSS protocol, making use of the underlying GSS-API network layer. rpc_gss_seccreate() allows an application to specify the type of security mechanism (for example, Kerberos v5), the type of service (for example, integrity checking), and the Quality of Protection (QOP) desired for transferring data.</p>														
<b>PARAMETERS</b>	<p>Information on RPCSEC_GSS data types for parameters may be found on the rpcsec_gss(3NSL) man page.</p> <table><tr><td><i>clnt</i></td><td>This is the RPC client handle. <i>clnt</i> may be obtained, for example, from <code>clnt_create()</code>.</td></tr><tr><td><i>principal</i></td><td>This is the identity of the server principal, specified in the form <i>service@host</i>, where <i>service</i> is the name of the service the client wishes to access and <i>host</i> is the fully qualified name of the host where the service resides — for example, <code>nfs@mymachine.eng.company.com</code>.</td></tr><tr><td><i>mechanism</i></td><td>This is an ASCII string which indicates which security mechanism to use with this data. Appropriate mechanisms may be found in the file <code>/etc/gss/mech</code>; additionally, <code>rpc_gss_get_mechanisms()</code> returns a list of supported security mechanisms (as null-terminated strings).</td></tr><tr><td><i>service_type</i></td><td>This sets the initial type of service for the session — privacy, integrity, authentication, or none.</td></tr><tr><td><i>qop</i></td><td>This is an ASCII string which sets the quality of protection (QOP) for the session. Appropriate values for this string may be found in the file <code>/etc/gss/qop</code>. Additionally, supported QOPs are returned (as null-terminated strings) by <code>rpc_gss_get_mech_info()</code>.</td></tr><tr><td><i>options_req</i></td><td>This structure contains options which are passed directly to the underlying GSS-API layer. If the caller specifies NULL for this parameter, defaults are used. (See NOTES, below.)</td></tr><tr><td><i>options_ret</i></td><td>These GSS-API options are returned to the caller. If the caller does not need to see these options, then it may specify NULL for this parameter. (See NOTES, below.)</td></tr></table>	<i>clnt</i>	This is the RPC client handle. <i>clnt</i> may be obtained, for example, from <code>clnt_create()</code> .	<i>principal</i>	This is the identity of the server principal, specified in the form <i>service@host</i> , where <i>service</i> is the name of the service the client wishes to access and <i>host</i> is the fully qualified name of the host where the service resides — for example, <code>nfs@mymachine.eng.company.com</code> .	<i>mechanism</i>	This is an ASCII string which indicates which security mechanism to use with this data. Appropriate mechanisms may be found in the file <code>/etc/gss/mech</code> ; additionally, <code>rpc_gss_get_mechanisms()</code> returns a list of supported security mechanisms (as null-terminated strings).	<i>service_type</i>	This sets the initial type of service for the session — privacy, integrity, authentication, or none.	<i>qop</i>	This is an ASCII string which sets the quality of protection (QOP) for the session. Appropriate values for this string may be found in the file <code>/etc/gss/qop</code> . Additionally, supported QOPs are returned (as null-terminated strings) by <code>rpc_gss_get_mech_info()</code> .	<i>options_req</i>	This structure contains options which are passed directly to the underlying GSS-API layer. If the caller specifies NULL for this parameter, defaults are used. (See NOTES, below.)	<i>options_ret</i>	These GSS-API options are returned to the caller. If the caller does not need to see these options, then it may specify NULL for this parameter. (See NOTES, below.)
<i>clnt</i>	This is the RPC client handle. <i>clnt</i> may be obtained, for example, from <code>clnt_create()</code> .														
<i>principal</i>	This is the identity of the server principal, specified in the form <i>service@host</i> , where <i>service</i> is the name of the service the client wishes to access and <i>host</i> is the fully qualified name of the host where the service resides — for example, <code>nfs@mymachine.eng.company.com</code> .														
<i>mechanism</i>	This is an ASCII string which indicates which security mechanism to use with this data. Appropriate mechanisms may be found in the file <code>/etc/gss/mech</code> ; additionally, <code>rpc_gss_get_mechanisms()</code> returns a list of supported security mechanisms (as null-terminated strings).														
<i>service_type</i>	This sets the initial type of service for the session — privacy, integrity, authentication, or none.														
<i>qop</i>	This is an ASCII string which sets the quality of protection (QOP) for the session. Appropriate values for this string may be found in the file <code>/etc/gss/qop</code> . Additionally, supported QOPs are returned (as null-terminated strings) by <code>rpc_gss_get_mech_info()</code> .														
<i>options_req</i>	This structure contains options which are passed directly to the underlying GSS-API layer. If the caller specifies NULL for this parameter, defaults are used. (See NOTES, below.)														
<i>options_ret</i>	These GSS-API options are returned to the caller. If the caller does not need to see these options, then it may specify NULL for this parameter. (See NOTES, below.)														

rpc\_gss\_seccreate(3NSL)

**RETURN VALUES** | `rpc_gss_seccreate()` returns a security context handle (an RPC authentication handle) of type AUTH. If `rpc_gss_seccreate()` cannot return successfully, the application can get an error number by calling `rpc_gss_get_error()`.

**FILES** | `/etc/gss/mech` | File containing valid security mechanisms  
`/etc/gss/qop` | File containing valid QOP values .

**ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe
Availability	SUNWrsg (32-bits)
	SUNWrsgx (64-bits)

**SEE ALSO** | `auth_destroy(3NSL)`, `rpc(3NSL)`, `rpc_gss_get_error(3NSL)`, `rpc_gss_get_mechanisms(3NSL)`, `rpcsec_gss(3NSL)`, `mech(4)`, `qop(4)`, `attributes(5)`

*ONC+ Developer's Guide*

Linn, J. *RFC 2743, Generic Security Service Application Program Interface Version 2, Update 1*. Network Working Group. January 2000.

**NOTES** | Contexts may be destroyed normally, with `auth_destroy()`. See `auth_destroy(3NSL)`

## rpc\_gss\_set\_callback(3NSL)

<b>NAME</b>	rpc_gss_set_callback – specify callback for context										
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpcsec_gss.h&gt;  bool_t <b>rpc_gss_set_callback</b>(struct rpc_gss_callback_t *cb);</pre>										
<b>DESCRIPTION</b>	<p>A server may want to specify a callback routine so that it knows when a context gets first used. This user-defined callback may be specified through the <code>rpc_gss_set_callback()</code> routine. The callback routine is invoked the first time a context is used for data exchanges, after the context is established for the specified program and version.</p> <p>The user-defined callback routine should take the following form:</p> <pre>bool_t callback(struct svc_req *req, gss_cred_id_t deleg,                gss_ctx_id_t gss_context, rpc_gss_lock_t *lock, void **cookie);</pre>										
<b>PARAMETERS</b>	<p><code>rpc_gss_set_callback()</code> takes one argument: a pointer to a <code>rpc_gss_callback_t</code> structure. This structure contains the RPC program and version number as well as a pointer to a user-defined <code>callback()</code> routine. (For a description of <code>rpc_gss_callback_t</code> and other <code>RPCSEC_GSS</code> data types, see the <code>rpcsec_gss(3NSL)</code> man page.)</p> <p>The user-defined <code>callback()</code> routine itself takes the following arguments:</p> <table><tr><td><i>req</i></td><td>Pointer to the received service request. <code>svc_req</code> is an RPC structure containing information on the context of an RPC invocation, such as program, version, and transport information.</td></tr><tr><td><i>deleg</i></td><td>Delegated credentials, if any. (See <code>NOTES</code>, below.)</td></tr><tr><td><i>gss_context</i></td><td>GSS context (allows server to do GSS operations on the context to test for acceptance criteria). See <code>NOTES</code>, below.</td></tr><tr><td><i>lock</i></td><td>This parameter is used to enforce a particular QOP and service for a session. This parameter points to a <code>RPCSEC_GSS rpc_gss_lock_t</code> structure. When the callback is invoked, the <code>rpc_gss_lock_t.locked</code> field is set to <code>TRUE</code>, thus locking the context. A locked context will reject all requests having different values for QOP or service than those specified by the <code>raw_cred</code> field of the <code>rpc_gss_lock_t</code> structure.</td></tr><tr><td><i>cookie</i></td><td>A four-byte quantity that an application may use in any manner it wants to — RPC does not interpret it. (For example, the cookie could be a pointer or index to a structure that represents a context initiator.) The cookie is returned, along with the caller's credentials, with each invocation of <code>rpc_gss_getcred()</code>.</td></tr></table>	<i>req</i>	Pointer to the received service request. <code>svc_req</code> is an RPC structure containing information on the context of an RPC invocation, such as program, version, and transport information.	<i>deleg</i>	Delegated credentials, if any. (See <code>NOTES</code> , below.)	<i>gss_context</i>	GSS context (allows server to do GSS operations on the context to test for acceptance criteria). See <code>NOTES</code> , below.	<i>lock</i>	This parameter is used to enforce a particular QOP and service for a session. This parameter points to a <code>RPCSEC_GSS rpc_gss_lock_t</code> structure. When the callback is invoked, the <code>rpc_gss_lock_t.locked</code> field is set to <code>TRUE</code> , thus locking the context. A locked context will reject all requests having different values for QOP or service than those specified by the <code>raw_cred</code> field of the <code>rpc_gss_lock_t</code> structure.	<i>cookie</i>	A four-byte quantity that an application may use in any manner it wants to — RPC does not interpret it. (For example, the cookie could be a pointer or index to a structure that represents a context initiator.) The cookie is returned, along with the caller's credentials, with each invocation of <code>rpc_gss_getcred()</code> .
<i>req</i>	Pointer to the received service request. <code>svc_req</code> is an RPC structure containing information on the context of an RPC invocation, such as program, version, and transport information.										
<i>deleg</i>	Delegated credentials, if any. (See <code>NOTES</code> , below.)										
<i>gss_context</i>	GSS context (allows server to do GSS operations on the context to test for acceptance criteria). See <code>NOTES</code> , below.										
<i>lock</i>	This parameter is used to enforce a particular QOP and service for a session. This parameter points to a <code>RPCSEC_GSS rpc_gss_lock_t</code> structure. When the callback is invoked, the <code>rpc_gss_lock_t.locked</code> field is set to <code>TRUE</code> , thus locking the context. A locked context will reject all requests having different values for QOP or service than those specified by the <code>raw_cred</code> field of the <code>rpc_gss_lock_t</code> structure.										
<i>cookie</i>	A four-byte quantity that an application may use in any manner it wants to — RPC does not interpret it. (For example, the cookie could be a pointer or index to a structure that represents a context initiator.) The cookie is returned, along with the caller's credentials, with each invocation of <code>rpc_gss_getcred()</code> .										
<b>RETURN VALUES</b>	<code>rpc_gss_set_callback()</code> returns <code>TRUE</code> if the use of the context is accepted; <code>false</code> otherwise.										
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:										



rpc\_gss\_set\_callback(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe
Availability	SUNWrsg (32-bit)
	SUNWrsgx (64-bit)

**SEE ALSO** `rpc(3NSL)`, `rpc_gss_getcred(3NSL)`, `rpcsec_gss(3NSL)`, `attributes(5)`

*ONC+ Developer's Guide*

Linn, J. *RFC 2078, Generic Security Service Application Program Interface, Version 2*. Network Working Group. January 1997.

**NOTES** If a server does not specify a callback, all incoming contexts will be accepted.

Because the GSS-API is not currently exposed, the *deleg* and *gss\_context* arguments are mentioned for informational purposes only, and the user-defined callback function may choose to do nothing with them.

## rpc\_gss\_set\_defaults(3NSL)

<b>NAME</b>	rpc_gss_set_defaults – change service, QOP for a session								
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpcsec_gss.h&gt;  bool_t <b>rpc_gss_set_defaults</b>(AUTH *auth, rpc_gss_service_t service,     char *qop);</pre>								
<b>DESCRIPTION</b>	rpc_gss_set_defaults() allows an application to change the service (privacy, integrity, authentication, or none) and Quality of Protection (QOP) for a transfer session. New values apply to the rest of the session (unless changed again).								
<b>PARAMETERS</b>	<p>Information on RPCSEC_GSS data types for parameters may be found on the rpcsec_gss(3NSL) man page.</p> <p><i>auth</i>                    An RPC authentication handle returned by rpc_gss_seccreate().</p> <p><i>service</i>                An enum of type rpc_gss_service_t, representing one of the following types of security service: authentication, privacy, integrity, or none.</p> <p><i>qop</i>                     A string representing Quality of Protection. Valid strings may be found in the file /etc/gss/qop or by using rpc_gss_get_mech_info().</p>								
<b>RETURN VALUES</b>	rpc_gss_set_svc_name() returns TRUE if it is successful; otherwise, use rpc_gss_get_error() to get the error associated with the failure.								
<b>FILES</b>	/etc/gss/qop                    File containing valid QOPs								
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:								
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> <tr> <td>Availability</td> <td>SUNWrsg (32-bit)</td> </tr> <tr> <td></td> <td>SUNWrsgx (64-bit)</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe	Availability	SUNWrsg (32-bit)		SUNWrsgx (64-bit)
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
MT-Level	MT-Safe								
Availability	SUNWrsg (32-bit)								
	SUNWrsgx (64-bit)								
<b>SEE ALSO</b>	<p>rpc(3NSL), rpc_gss_get_mech_info(3NSL), rpcsec_gss(3NSL), qop(4), attributes(5)</p> <p><i>ONC+ Developer's Guide</i></p> <p>Linn, J. <i>RFC 2078, Generic Security Service Application Program Interface, Version 2.</i> Network Working Group. January 1997.</p>								

<b>NAME</b>	rpc_gss_set_svc_name – send a principal name to a server								
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpcsec_gss.h&gt;  bool_t <b>rpc_gss_set_svc_name</b>(char *principal, char *mechanism, u_int     req_time, u_int program, u_int version);</pre>								
<b>DESCRIPTION</b>	rpc_gss_set_svc_name() sets the name of a principal the server is to represent. If a server is going to act as more than one principal, this procedure can be invoked for every such principal.								
<b>PARAMETERS</b>	<p>Information on RPCSEC_GSS data types for parameters may be found on the rpcsec_gss(3NSL) man page.</p> <p><i>principal</i>            An ASCII string representing the server's principal name, given in the form of <i>service@host</i>.</p> <p><i>mech</i>                    An ASCII string representing the security mechanism in use. Valid strings may be found in the <i>/etc/gss/mech</i> file, or by using <i>rpc_gss_get_mechanisms()</i>.</p> <p><i>req_time</i>                The time, in seconds, for which a credential should be valid. Note that the <i>req_time</i> is a hint to the underlying mechanism. The actual time that the credential will remain valid is mechanism dependent. In the case of kerberos the actual time will be GSS_C_INDEFINITE.</p> <p><i>program</i>                The RPC program number for this service.</p> <p><i>version</i>                 The RPC version number for this service.</p>								
<b>RETURN VALUES</b>	rpc_gss_set_svc_name() returns TRUE if it is successful; otherwise, use <i>rpc_gss_get_error()</i> to get the error associated with the failure.								
<b>FILES</b>	<i>/etc/gss/mech</i> File containing valid security mechanisms								
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:								
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> <tr> <td>Availability</td> <td>SUNWrsg (32-bit)</td> </tr> <tr> <td></td> <td>SUNWrsgx (64-bit)</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe	Availability	SUNWrsg (32-bit)		SUNWrsgx (64-bit)
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
MT-Level	MT-Safe								
Availability	SUNWrsg (32-bit)								
	SUNWrsgx (64-bit)								
<b>SEE ALSO</b>	<p>rpc(3NSL), <i>rpc_gss_get_mechanisms(3NSL)</i>, <i>rpc_gss_get_principal_name(3NSL)</i>, <i>rpcsec_gss(3NSL)</i>, <i>mech(4)</i>, <i>attributes(5)</i></p> <p><i>ONC+ Developer's Guide</i></p>								

rpc\_gss\_set\_svc\_name(3NSL)

Linn, J. *RFC 2078, Generic Security Service Application Program Interface, Version 2.*  
Network Working Group. January 1997.

<b>NAME</b>	rpc_rac, rac_drop, rac_poll, rac_recv, rac_send – remote asynchronous calls						
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lrac -lnsl [ library ... ] #include &lt;rpc/rpc.h&gt; #include &lt;rpc/rac.h&gt;  void <b>rac_drop</b>(CLIENT *cl, void *h);  enum clnt_stat <b>rac_poll</b>(CLIENT *cl, void *h);  enum clnt_stat <b>rac_recv</b>(CLIENT *cl, void *h);  void *<b>rac_send</b>(CLIENT *cl, rpcproc_t proc, xdrproc_t xargs, void                *argsp, xdrproc_t xresults, void *resultsp, struct timeval timeout);</pre>						
<b>DESCRIPTION</b>	<p>The remote asynchronous calls (RAC) package is a special interface to the RPC library that allows messages to be sent using the RPC protocol without blocking during the time between when the message is sent and the reply is received. To RPC servers, RAC messages are indistinguishable from RPC messages.</p> <p>A client establishes an RPC session in the usual way (see <code>rpc_clnt_create(3NSL)</code>). A RAC message is sent using <code>rac_send()</code>. This routine returns immediately, allowing the client to conduct other processing. When the client wants to determine whether the returned value from the call has been received, <code>rac_poll()</code> is used. <code>rac_recv()</code> is used to collect the returned value; it can also be used to block while waiting for the returned value to arrive. <code>rac_drop()</code> is used to inform the RPC library that the client is no longer interested in the results of a particular RAC message.</p> <p><code>rac_drop()</code>      <code>rac_drop()</code> should be called when the user is no longer interested in the result of a <code>rac_send()</code> currently in progress. No message to the server is generated by this call, but any subsequent reply received for this handle will be silently dropped. It also frees any space occupied by the asynchronous call handle <i>h</i>.</p> <p>After a call to <code>rac_drop()</code> the handle referred to by <i>h</i> is invalid. It may no longer be used in any asynchronous operation.</p> <p><code>rac_poll()</code>      <code>rac_poll()</code> returns the status of the call currently in progress on the &lt;CLIENT, asynchronous handle&gt; tuple referred to by <i>cl</i> and <i>h</i>.</p> <p><code>rac_poll()</code> return values are:</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="padding-right: 2em;">RPC_SUCCESS</td> <td>A reply has been received and is available for reading by <code>rac_recv()</code>.</td> </tr> <tr> <td>RPC_INPROGRESS</td> <td>No reply has been received. The call referred to by the given handle has not yet timed out.</td> </tr> <tr> <td>RPC_TIMEDOUT</td> <td>No reply has been received. The call referred to by the given handle</td> </tr> </table>	RPC_SUCCESS	A reply has been received and is available for reading by <code>rac_recv()</code> .	RPC_INPROGRESS	No reply has been received. The call referred to by the given handle has not yet timed out.	RPC_TIMEDOUT	No reply has been received. The call referred to by the given handle
RPC_SUCCESS	A reply has been received and is available for reading by <code>rac_recv()</code> .						
RPC_INPROGRESS	No reply has been received. The call referred to by the given handle has not yet timed out.						
RPC_TIMEDOUT	No reply has been received. The call referred to by the given handle						

## rpc\_rac(3RAC)

	has exceeded the maximum timeout value specified in <code>rac_send()</code> .
<code>RPC_STALERACHANDLE</code>	Either the handle referred to by <i>h</i> is invalid or no call is currently in progress for the given <CLIENT, asynchronous handle> tuple.
<code>RPC_CANTRECV</code>	Either the file descriptor associated with the given CLIENT handle is bad, or an error occurred while attempting to receive a packet.
<code>RPC_SYSTEMERROR</code>	Space could not be allocated to receive a packet.

On unreliable transports, a call to `rac_poll()` will trigger a retransmission when necessary (that is, if a `rac_send()` is in progress, no reply has been received, the per-call timeout has expired, and the total timeout has not yet expired).

The return value for `rac_poll()` is independent of the RPC return value in the reply packet. Although a combination of `clnt_control()`'s `CLGET_FD` request and `poll(2)` may be used to extract the proper file descriptor and poll for packets, `rac_poll()` is still useful since it will determine whether a reply is available for a specific <CLIENT, asynchronous handle> tuple.

`rac_recv()` retrieves the results of a previous asynchronous RPC call, placing them in the buffer indicated in the `rac_send()` call and using the XDR decode function supplied there. It depends on the application to have ensured that a reply is present (using `rac_poll()`). If `rac_recv()` is called before a reply has been received, it will block awaiting a reply.

All errors normally returned by the RPC client call functions may be returned here. In addition:

<code>RPC_STALERACHANDLE</code>	Either the handle referred to by <i>h</i> is invalid or no call is currently in progress for the given <CLIENT, asynchronous handle> tuple.
	Additionally, if a packet is present and its status is not <code>RPC_SUCCESS</code> , it is possible that the client credentials need refreshing. In this case, <code>RPC_AUTHERROR</code> is returned and the client should

rpc\_rac(3RAC)

attempt to resend the call.

When a reply has been received, `rac_recv()` will invoke the XDR decode procedure specified in the `rac_send()` call. After a call to `rac_recv()`, the handle referred to by `h` is invalid. It may no longer be used in any asynchronous operation.

`rac_send()` `rac_send()` initiates (sends to the server) an RPC call to the specified procedure. It does not await a reply from the server. `argsp` is the address of the procedure's arguments, `resultsp` is the address in which to place the results, `xargs` and `xresults` are XDR functions used to encode and decode respectively. Note: `resultsp` must be a valid pointer when `rac_recv()` is called. `timeout` should contain the total amount of time the application is willing to wait for a reply.

Upon success, an opaque handle, known as the asynchronous handle, is returned. This handle is to be used in subsequent asynchronous calls to poll for the status of the call (`rac_poll()`), receive the returned results of the call (`rac_recv()`), or cancel the call (`rac_drop()`).

On failure, `(void *) 0` is returned.

In case of failure, the application may retrieve the RPC failure code by calling `clnt_geterr()` immediately after a `rac_send()` failure (see `rpc(3NSL)`). Possible errors include both transient problems (such as transport failures) and permanent ones (such as XDR encoding failures).

Multiple `rac_sends` on the same client handle are permitted, but may introduce unpredictable perturbations to the current timeout and retry model used by the RPC library.

The interface imposes a limit on the amount of time a call may be in progress before it is considered to have failed. This method was chosen over limitations on the number of retries because of a desire for transport independence.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO** `poll(2)`, `rpc(3NSL)`, `rpc_clnt_create(3NSL)`, `rpc_clnt_calls(3NSL)`, `xdr(3NSL)`, `attributes(5)`

rpc\_rac(3RAC)

**WARNINGS** The RAC interface is not the recommended interface for having multiple RPC requests outstanding. The preferred method of accomplishing this in the Solaris environment is to use synchronous RPC calls with threads. The RAC interface is provided as a service to developers interested in porting RPC applications to Solaris 2.0. Use of this interface will degrade the performance of normal synchronous RPC calls (see `rpc_clnt_calls(3NSL)`). For these reasons, use of this interface is disparaged.

The library `librac` must be linked before `libns1` to use RAC. If the libraries are not linked in the correct order, then the results are indeterminate.

**NOTES** These interfaces are unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.



<b>NAME</b>	rpcsec_gss – security flavor incorporating GSS-API protections
<b>SYNOPSIS</b>	<pre>cc [ flag... ] file... -lnsl [ library... ] #include &lt;rpc/rpcsec_gss.h&gt;</pre>
<b>DESCRIPTION</b>	<p>RPCSEC_GSS is a security flavor which sits "on top" of the GSS-API (Generic Security Service API) for network transmissions. Applications using RPCSEC_GSS can take advantage of GSS-API security features; moreover, they can use any security mechanism (such as RSA public key or Kerberos) that works with the GSS-API.</p> <p>The GSS-API offers two security services beyond the traditional authentication services (AUTH_DH, AUTH_SYS, and AUTH_KERB): integrity and privacy. With integrity, the system uses cryptographic checksumming to ensure the authenticity of a message (authenticity of originator, recipient, and data); privacy provides additional security by encrypting data. Applications using RPCSEC_GSS specify which service they wish to use. Type of security service is mechanism-independent.</p> <p>Before exchanging data with a peer, an application must establish a context for the exchange. RPCSEC_GSS provides a single function for this purpose, <code>rpc_gss_seccreate()</code>, which allows the application to specify the security mechanism, Quality of Protection (QOP), and type of service at context creation. (The QOP parameter sets the cryptographic algorithms to be used with integrity or privacy, and is mechanism-dependent.) Once a context is established, applications can reset the QOP and type of service for each data unit exchanged, if desired.</p> <p>Valid mechanisms and QOPs may be obtained from configuration files or from the name service. Each mechanism has a default QOP.</p> <p>Contexts are destroyed with the usual RPC <code>auth_destroy()</code> call.</p>
<b>Data Structures</b>	<p>Some of the data structures used by the RPCSEC_GSS package are shown below.</p> <p><code>rpc_gss_service_t</code></p> <p>This enum defines the types of security services the context may have. <code>rpc_gss_seccreate()</code> takes this as one argument when setting the service type for a session.</p> <pre>typedef enum {     rpc_gss_svc_default = 0,     rpc_gss_svc_none = 1,     rpc_gss_svc_integrity = 2,     rpc_gss_svc_privacy = 3 } rpc_gss_service_t ;</pre> <p><code>rpc_gss_options_req_t</code></p> <p>Structure containing options passed directly through to the GSS-API. <code>rpc_gss_seccreate()</code> takes this as an argument when creating a context.</p> <pre>typedef struct {     int req_flags; /*GSS request bits */     int time_req; /*requested credential lifetime */</pre>

## rpcsec\_gss(3NSL)

```
    gss_cred_id_t my_cred; /*GSS credential struct*/
    gss_channel_bindings_t;
    input_channel_bindings;
} rpc_gss_options_req_t ;
```

### rpc\_gss\_OID

This data type is used by in-kernel RPC routines, and thus is mentioned here for informational purposes only.

```
typedef struct {
    u_int    length;
    void    *elements
} *rpc_gss_OID;
```

### rpc\_gss\_options\_ret\_t

Structure containing GSS-API options returned to the calling function, `rpc_gss_seccreate()`. `MAX_GSS_MECH` is defined as 128.

```
typedef struct {
    int        major_status;
    int        minor_status;
    u_int      rpcsec_version           /*vers. of RPCSEC_GSS */
    int        ret_flags
    int        time_req
    gss_ctx_id_t gss_context;
    char       actual_mechanism[MAX_GSS_MECH]; /*mechanism used*/
} rpc_gss_options_ret_t;
```

### rpc\_gss\_principal\_t

The (mechanism-dependent, opaque) client principal type. Used as an argument to the `rpc_gss_get_principal_name()` function, and in the `gsscred` table. Also referenced by the `rpc_gss_rawcred_t` structure for raw credentials (see below).

```
typedef struct {
    int len;
    char name[1];
} *rpc_gss_principal_t;
```

### rpc\_gss\_rawcred\_t

Structure for raw credentials. Used by `rpc_gss_getcred()` and `rpc_gss_set_callback()`.

```
typedef struct {
    u_int        version;           /*RPC version # */
    char         *mechanism;        /*security mechanism*/
    char         *qop;              /*Quality of Protection*/
    rpc_gss_principal_t client_principal; /*client name*/
    char         *svc_principal;    /*server name*/
    rpc_gss_service_t service;     /*service (integrity, etc.)*/
} rpc_gss_rawcred_t;
```

rpc\_gss\_ucred\_t

Structure for UNIX credentials. Used by `rpc_gss_getcred()` as an alternative to `rpc_gss_rawcred_t`.

```
typedef struct {
    uid_t  uid;      /*user ID*/
    gid_t  gid;      /*group ID*/
    short  gidlen;
    git_t  *gidlist; /*list of groups*/
} rpc_gss_ucred_t;
```

rpc\_gss\_callback\_t

Callback structure used by `rpc_gss_set_callback()`.

```
typedef struct {
    u_int  program; /*RPC program #*/
    u_int  version; /*RPC version #*/
    bool_t (*callback)(); /*user-defined callback routine*/
} rpc_gss_callback_t;
```

rpc\_gss\_lock\_t

Structure used by a callback routine to enforce a particular QOP and service for a session. The `locked` field is normally set to `FALSE`; the server sets it to `TRUE` in order to lock the session. (A locked context will reject all requests having different QOP and service values than those found in the `raw_cred` structure.) For more information, see the `rpc_gss_set_callback(3NSL)` man page.

```
typedef struct {
    bool_t      locked;
    rpc_gss_rawcred_t *raw_cred;
} rpc_gss_lock_t;
```

rpc\_gss\_error\_t

Structure used by `rpc_gss_get_error()` to fetch an error code when a `RPCSEC_GSS` routine fails.

```
typedef struct {
    int  rpc_gss_error;
    int  system_error; /*same as errno*/
} rpc_gss_error_t;
```

## Index to Routines

The following lists `RPCSEC_GSS` routines and the manual reference pages on which they are described. An (S) indicates it is a server-side function:

Routine (Manual Page)

Description

`rpc_gss_seccreate(3NSL)`

Create a secure `RPCSEC_GSS` context

## rpcsec\_gss(3NSL)

`rpc_gss_set_defaults(3NSL)`  
Switch service, QOP for a session

`rpc_gss_max_data_length(3NSL)`  
Get maximum data length allowed by transport

`rpc_gss_set_svc_name(3NSL)`  
Set server's principal name (S)

`rpc_gss_getcred(3NSL)`  
Get credentials of caller (S)

`rpc_gss_set_callback(3NSL)`  
Specify callback to see context use (S)

`rpc_gss_get_principal_name(3NSL)`  
Get client principal name (S)

`rpc_gss_svc_max_data_length(3NSL)`  
Get maximum data length allowed by transport (S)

`rpc_gss_get_error(3NSL)`  
Get error number

`rpc_gss_get_mechanisms(3NSL)`  
Get valid mechanism strings

`rpc_gss_get_mech_info(3NSL)`  
Get valid QOP strings, current service

`rpc_gss_get_versions(3NSL)`  
Get supported RPCSEC\_GSS versions

`rpc_gss_is_installed(3NSL)`  
Checks if a mechanism is installed

`rpc_gss_mech_to_oid(3NSL)`  
Maps ASCII mechanism to OID representation

`rpc_gss_qop_to_num(3NSL)`  
Maps ASCII QOP, mechanism to u\_int number

**Utilities** The `gsscred` utility manages the `gsscred` table, which contains mappings of principal names between network and local credentials. See `gsscred(1M)`.

**FILES** `/etc/gss/mech` List of installed mechanisms  
`/etc/gss/qop` List of valid QOPs

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

rpcsec\_gss(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWrsg (32-bit)
	SUNWrsgx (64-bit)

**SEE ALSO** gsscred(1M), rpc(3NSL), rpc\_clnt\_auth(3NSL), xdr(3NSL), attributes(5), environ(5)

*ONC+ Developer's Guide*

Linn, J. *RFC 2743, Generic Security Service Application Program Interface Version 2, Update 1*. Network Working Group. January 2000.

## rpc\_soc(3NSL)

<b>NAME</b>	rpc_soc, authdes_create, authunix_create, authunix_create_default, callrpc, clnt_broadcast, clntraw_create, clnttcp_create, clntudp_bufcreate, clntudp_create, get_myaddress, getrpcport, pmap_getmaps, pmap_getport, pmap_rmtcall, pmap_set, pmap_unset, registerrpc, svc_fds, svc_getcaller, svc_getreq, svc_register, svc_unregister, svcd_create, svcraw_create, svctcp_create, svcudp_bufcreate, svcudp_create, xdr_authunix_parms – obsolete library routines for RPC
<b>SYNOPSIS</b>	<pre>#define PORTMAP #include &lt;rpc/rpc.h&gt;  AUTH *authdes_create(char *name, uint_t window, struct sockaddr_in     *syncaddr, des_block *ckey);  AUTH *authunix_create(char *host, uid_t uid, gid_t gid, int grouplen,     gid_t *gidlistp);  AUTH *authunix_create_default(void);  callrpc(char *host, rpcprog_t prognum, rpcvers_t versnum, rpcproc_t     procnum, xdrproc_t inproc, char *in, xdrproc_t outproc, char *out);  enum clnt_stat clnt_broadcast(rpcprog_t prognum, rpcvers_t     versnum, rpcproc_t procnum, xdrproc_t inproc, char *in, xdrproc_t     outproc, char *out, resultproc_t teachresult);  CLIENT *clntraw_create(rpcproc_t procnum, rpcvers_t versnum);  CLIENT *clnttcp_create(struct sockaddr_in *addr, rpcprog_t     prognum, rpcvers_t versnum, int *fdp, uint_t sendz, uint_t recvsz);  CLIENT *clntudp_bufcreate(struct sockaddr_in *addr, rpcprog_t     prognum, rpcvers_t versnum, struct timeval wait, int *fdp, uint_t     sendz, uint_t recvsz);  CLIENT *clntudp_create(struct sockaddr_in *addr, rpcprog_t     prognum, struct timeval wait, int *fdp);  void get_myaddress(struct sockaddr_in *addr);  ushort getrpcport(char *host, rpcprog_t prognum, rpcvers_t versnum,     rpcprot_t proto);  struct pmaplist *pmap_getmaps(struct sockaddr_in *addr);  ushort pmap_getport(struct sockaddr_in *addr, rpcprog_t prognum,     rpcvers_t versnum, rpcprot_t protocol);  enum clnt_stat pmap_rmtcall(struct sockaddr_in *addr, rpcprog_t     prognum, rpcvers_t versnum, rpcproc_t prognum, caddr_t in,     xdrproc_t inproc, caddr_t out, xdrproc_t outproc, struct     timeval tout, rpcport_t *portp);  bool_t pmap_set(rpcprog_t prognum, rpcvers_t versnum, rpcprot_t     protocol, u_short port);  bool_t pmap_unset(rpcprog_t prognum, rpcvers_t versnum);</pre>

```

int svc_fds;
struct sockaddr_in *svc_getcaller(SVCXPRT *xprt);
void svc_getreq(int rdfs);
SVCXPRT *svcfid_create(int fd, uint_t sendsz, uint_t recvsz);
SVCXPRT *svcrow_create(void);
SVCXPRT *svctcp_create(int fd, uint_t sendsz, uint_t recvsz);
SVCXPRT *svcudp_bufcreate(int fd, uint_t sendsz, uint_t recvsz);
SVCXPRT *svcudp_create(int fd);
registerrpc(rpcprog_t prognum, rpcvers_t versnum, rpcproc_t procnum,
            char *(*procname)(), xdrproc_t inproc, xdrproc_t outproc);
bool_t svc_register(SVCXPRT *xprt, rpcprog_t prognum, rpcvers_t
                   versnum, void (*dispatch)(), int protocol);
void svc_unregister(rpcprog_t prognum, rpcvers_t versnum);
bool_t xdr_authunix_parms(XDR *xdrs, struct authunix_parms *supp);

```

**DESCRIPTION** RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

The routines described in this manual page have been superseded by other routines. The preferred routine is given after the description of the routine. New programs should use the preferred routines, as support for the older interfaces may be dropped in future releases.

**File Descriptors** Transport independent RPC uses TLI as its transport interface instead of sockets.

Some of the routines described in this section (such as `clnttcp_create()`) take a pointer to a file descriptor as one of the parameters. If the user wants the file descriptor to be a socket, then the application will have to be linked with both `librpcsoc` and `libnsl`. If the user passed `RPC_ANYSOCK` as the file descriptor, and the application is linked with `libnsl` only, then the routine will return a TLI file descriptor and not a socket.

**Routines** The following routines require that the header `<rpc/rpc.h>` be included. The symbol `PORTMAP` should be defined so that the appropriate function declarations for the old interfaces are included through the header files.

```
authdes_create()
```

`authdes_create()` is the first of two routines which interface to the RPC secure authentication system, known as DES authentication. The second is `authdes_getucred()`, below. Note: the keyserver daemon `keyserv(1M)` must be running for the DES authentication system to work.

## rpc\_soc(3NSL)

`authdes_create()`, used on the client side, returns an authentication handle that will enable the use of the secure authentication system. The first parameter *name* is the network name, or *netname*, of the owner of the server process. This field usually represents a hostname derived from the utility routine `host2netname()`, but could also represent a user name using `user2netname()`. See `secure_rpc(3NSL)`. The second field is window on the validity of the client credential, given in seconds. A small window is more secure than a large one, but choosing too small of a window will increase the frequency of resynchronizations because of clock drift. The third parameter *syncaddr* is optional. If it is `NULL`, then the authentication system will assume that the local clock is always in sync with the server's clock, and will not attempt resynchronizations. If an address is supplied, however, then the system will use the address for consulting the remote time service whenever resynchronization is required. This parameter is usually the address of the RPC server itself. The final parameter *key* is also optional. If it is `NULL`, then the authentication system will generate a random DES key to be used for the encryption of credentials. If it is supplied, however, then it will be used instead.

This routine exists for backward compatibility only, and it is made obsolete by `authdes_seccreate()`. See `secure_rpc(3NSL)`.

### `authunix_create()`

Create and return an RPC authentication handle that contains .UX authentication information. The parameter *host* is the name of the machine on which the information was created; *uid* is the user's user ID; *gid* is the user's current group ID; *grouplen* and *gidlistp* refer to a counted array of groups to which the user belongs.

It is not very difficult to impersonate a user.

This routine exists for backward compatibility only, and it is made obsolete by `authsys_create()`. See `rpc_clnt_auth(3NSL)`.

### `authunix_create_default()`

Call `authunix_create()` with the appropriate parameters.

This routine exists for backward compatibility only, and it is made obsolete by `authsys_create_default()`. See `rpc_clnt_auth(3NSL)`.

### `callrpc()`

Call the remote procedure associated with *prognum*, *versnum*, and *procnum* on the machine, *host*. The parameter *inproc* is used to encode the procedure's parameters, and *outproc* is used to decode the procedure's results; *in* is the address of the procedure's argument, and *out* is the address of where to place the result(s). This routine returns 0 if it succeeds, or the value of enum `clnt_stat` cast to an integer if it fails. The routine `clnt_perrno()` is handy for translating failure statuses into messages. See `rpc_clnt_calls(3NSL)`.

You do not have control of timeouts or authentication using this routine. This routine exists for backward compatibility only, and is made obsolete by `rpc_call()`. See `rpc_clnt_calls(3NSL)`.



`clnt_stat_clnt_broadcast()`

Like `callrpc()`, except the call message is broadcast to all locally connected broadcast nets. Each time the caller receives a response, this routine calls `eachresult()`, whose form is:

```
eachresult(char *out, struct sockaddr_in *addr);
```

where `out` is the same as `out` passed to `clnt_broadcast()`, except that the remote procedure's output is decoded there; `addr` points to the address of the machine that sent the results. If `eachresult()` returns 0, `clnt_broadcast()` waits for more replies; otherwise it returns with appropriate status. If `eachresult()` is NULL, `clnt_broadcast()` returns without waiting for any replies.

Broadcast packets are limited in size to the maximum transfer unit of the transports involved. For Ethernet, the caller's argument size is approximately 1500 bytes. Since the call message is sent to all connected networks, it may potentially lead to broadcast storms. `clnt_broadcast()` uses SB AUTH\_SYS credentials by default. See `rpc_clnt_auth(3NSL)`. This routine exists for backward compatibility only, and is made obsolete by `rpc_broadcast()`. See `rpc_clnt_calls(3NSL)`.

`clntraw_create()`

This routine creates an internal, memory-based RPC client for the remote program `prognum`, version `versnum`. The transport used to pass messages to the service is actually a buffer within the process's address space, so the corresponding RPC server should live in the same address space. See `svcrw_create()`. This allows simulation of RPC and acquisition of RPC overheads, such as round trip times, without any kernel interference. This routine returns NULL if it fails.

This routine exists for backward compatibility only. It has the same functionality as `clnt_raw_create()`. See `rpc_clnt_create(3NSL)`, which obsoletes it.

`clnttcp_create()`

This routine creates an RPC client for the remote program `prognum`, version `versnum`; the client uses TCP/IP as a transport. The remote program is located at Internet address `addr`. If `addr->sin_port` is 0, then it is set to the actual port that the remote program is listening on. The remote `rpcbind` service is consulted for this information. The parameter `*fdp` is a file descriptor, which may be open and bound; if it is `RPC_ANYSOCK`, then this routine opens a new one and sets `*fdp`. Refer to the File Descriptor section for more information. Since TCP-based RPC uses buffered I/O, the user may specify the size of the send and receive buffers with the parameters `sendsz` and `recvsz`. Values of 0 choose suitable defaults. This routine returns NULL if it fails.

This routine exists for backward compatibility only. `clnt_create()`, `clnt_tli_create()`, or `clnt_vc_create()` should be used instead. See `rpc_clnt_create(3NSL)`.

`clntudp_bufcreate()`

Create a client handle for the remote program `prognum`, on `versnum`; the client uses UDP/IP as the transport. The remote program is located at the Internet address `addr`. If `addr->sin_port` is 0, it is set to port on which the remote program is listening

## rpc\_soc(3NSL)

on (the remote `rpcbind` service is consulted for this information). The parameter `*fdp` is a file descriptor, which may be open and bound. If it is `RPC_ANYSOCK`, then this routine opens a new one and sets `*fdp`. Refer to the `File Descriptor` section for more information. The UDP transport resends the call message in intervals of wait time until a response is received or until the call times out. The total time for the call to time out is specified by `clnt_call()`. See `rpc_clnt_calls(3NSL)`. If successful it returns a client handle, otherwise it returns `NULL`. The error can be printed using the `clnt_pcreateerror()` routine. See `rpc_clnt_create(3NSL)`.

The user can specify the maximum packet size for sending and receiving by using `sendsz` and `recvsz` arguments for UDP-based RPC messages.

If `addr->sin_port` is 0 and the requested version number `versnum` is not registered with the remote portmap service, it returns a handle if at least a version number for the given program number is registered. The version mismatch is discovered by a `clnt_call()` later (see `rpc_clnt_calls(3NSL)`).

This routine exists for backward compatibility only. `clnt_tli_create()` or `clnt_dg_create()` should be used instead. See `rpc_clnt_create(3NSL)`.

### `clntudp_create()`

This routine creates an RPC client handle for the remote program `prognum`, version `versnum`; the client uses UDP/IP as a transport. The remote program is located at Internet address `addr`. If `addr->sin_port` is 0, then it is set to actual port that the remote program is listening on. The remote `rpcbind` service is consulted for this information. The parameter `*fdp` is a file descriptor, which may be open and bound; if it is `RPC_ANYSOCK`, then this routine opens a new one and sets `*fdp`. Refer to the `File Descriptor` section for more information. The UDP transport resends the call message in intervals of wait time until a response is received or until the call times out. The total time for the call to time out is specified by `clnt_call()`. See `rpc_clnt_calls(3NSL)`. `clntudp_create()` returns a client handle on success, otherwise it returns `NULL`. The error can be printed using the `clnt_pcreateerror()` routine. See `rpc_clnt_create(3NSL)`.

Since UDP-based RPC messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take large arguments or return huge results.

This routine exists for backward compatibility only. `clnt_create()`, `clnt_tli_create()`, or `clnt_dg_create()` should be used instead. See `rpc_clnt_create(3NSL)`.

### `get_myaddress()`

Places the local system's IP address into `*addr`, without consulting the library routines that deal with `/etc/hosts`. The port number is always set to `htons(PMAPPORT)`.

This routine is only intended for use with the RPC library. It returns the local system's address in a form compatible with the RPC library, and should not be taken as the system's actual IP address. In fact, the *\*addr* buffer's host address part is actually zeroed. This address may have only local significance and should not be assumed to be an address that can be used to connect to the local system by remote systems or processes.

This routine remains for backward compatibility only. The routine `netdir_getbyname()` should be used with the name `HOST_SELF` to retrieve the local system's network address as a *netbuf* structure. See `netdir(3NSL)`.

#### `getrpcport()`

`getrpcport()` returns the port number for the version *versnum* of the RPC program *prognum* running on *host* and using protocol *proto*. `getrpcport()` returns 0 if the RPC system failed to contact the remote portmap service, the program associated with *prognum* is not registered, or there is no mapping between the program and a port.

This routine exists for backward compatibility only. Enhanced functionality is provided by `rpcb_getaddr()`. See `rpcbind(3NSL)`.

#### `pmaplist()`

A user interface to the portmap service, which returns a list of the current RPC program-to-port mappings on the host located at IP address *addr*. This routine can return NULL. The command `'rpcinfo -p'` uses this routine.

This routine exists for backward compatibility only, enhanced functionality is provided by `rpcb_getmaps()`. See `rpcbind(3NSL)`.

#### `pmap_getport()`

A user interface to the portmap service, which returns the port number on which waits a service that supports program *prognum*, version *versnum*, and speaks the transport protocol associated with *protocol*. The value of *protocol* is most likely `IPPROTO_UDP` or `IPPROTO_TCP`. A return value of 0 means that the mapping does not exist or that the RPC system failed to contact the remote portmap service. In the latter case, the global variable `rpc_createerr` contains the RPC status.

This routine exists for backward compatibility only, enhanced functionality is provided by `rpcb_getaddr()`. See `rpcbind(3NSL)`.

#### `pmap_rmtcall()`

Request that the portmap on the host at IP address *\*addr* make an RPC on the behalf of the caller to a procedure on that host. *\*portp* is modified to the program's port number if the procedure succeeds. The definitions of other parameters are discussed in `callrpc()` and `clnt_call()`. See `rpc_clnt_calls(3NSL)`.

This procedure is only available for the UDP transport.

If the requested remote procedure is not registered with the remote portmap then no error response is returned and the call times out. Also, no authentication is done.

## rpc\_soc(3NSL)

This routine exists for backward compatibility only, enhanced functionality is provided by `rpcb_rmtcall()`. See `rpcbind(3NSL)`.

### `pmap_set()`

A user interface to the `portmap` service, that establishes a mapping between the triple [*prognum*, *versnum*, *protocol*] and *port* on the machine's `portmap` service. The value of *protocol* may be `IPPROTO_UDP` or `IPPROTO_TCP`. Formerly, the routine failed if the requested *port* was found to be in use. Now, the routine only fails if it finds that *port* is still bound. If *port* is not bound, the routine completes the requested registration. This routine returns 1 if it succeeds, 0 otherwise. Automatically done by `svc_register()`.

This routine exists for backward compatibility only, enhanced functionality is provided by `rpcb_set()`. See `rpcbind(3NSL)`.

### `pmap_unset()`

A user interface to the `portmap` service, which destroys all mapping between the triple [*prognum*, *versnum*, *all-protocols*] and *port* on the machine's `portmap` service. This routine returns one if it succeeds, 0 otherwise.

This routine exists for backward compatibility only, enhanced functionality is provided by `rpcb_unset()`. See `rpcbind(3NSL)`.

### `svc_fds()`

A global variable reflecting the RPC service side's read file descriptor bit mask; it is suitable as a parameter to the `select()` call. This is only of interest if a service implementor does not call `svc_run()`, but rather does his own asynchronous event processing. This variable is read-only, yet it may change after calls to `svc_getreq()` or any creation routines. Do not pass its address to `select()`! Similar to `svc_fdset`, but limited to 32 descriptors.

This interface is made obsolete by `svc_fdset`. See `rpc_svc_calls(3NSL)`.

### `svc_getcaller()`

This routine returns the network address, represented as a `struct sockaddr_in`, of the caller of a procedure associated with the RPC service transport handle, *xprt*.

This routine exists for backward compatibility only, and is obsolete. The preferred interface is `svc_getrpccaller()`. See `rpc_svc_reg(3NSL)`, which returns the address as a `struct netbuf`.

### `svc_getreq()`

This routine is only of interest if a service implementor does not call `svc_run()`, but instead implements custom asynchronous event processing. It is called when the `select()` call has determined that an RPC request has arrived on some RPC file descriptors; *rdfds* is the resultant read file descriptor bit mask. The routine returns when all file descriptors associated with the value of *rdfds* have been serviced. This routine is similar to `svc_getreqset()` but is limited to 32 descriptors.

This interface is made obsolete by `svc_getreqset()`

`svcfcreate()`

Create a service on top of any open and bound descriptor. Typically, this descriptor is a connected file descriptor for a stream protocol. Refer to the `File Descriptor` section for more information. `sendsz` and `recvsz` indicate sizes for the send and receive buffers. If they are 0, a reasonable default is chosen.

This interface is made obsolete by `svc_fd_create()` (see `rpc_svc_create(3NSL)`).

`svccraw_create()`

This routine creates an internal, memory-based RPC service transport, to which it returns a pointer. The transport is really a buffer within the process's address space, so the corresponding RPC client should live in the same address space; see `clntraw_create()`. This routine allows simulation of RPC and acquisition of RPC overheads (such as round trip times), without any kernel interference. This routine returns NULL if it fails.

This routine exists for backward compatibility only, and has the same functionality of `svc_raw_create()`. See `rpc_svc_create(3NSL)`, which obsoletes it.

`svctcp_create()`

This routine creates a TCP/IP-based RPC service transport, to which it returns a pointer. The transport is associated with the file descriptor `fd`, which may be `RPC_ANYSOCK`, in which case a new file descriptor is created. If the file descriptor is not bound to a local TCP port, then this routine binds it to an arbitrary port. Refer to the `File Descriptor` section for more information. Upon completion, `xprt->xp_fd` is the transport's file descriptor, and `xprt->xp_port` is the transport's port number. This routine returns NULL if it fails. Since TCP-based RPC uses buffered I/O, users may specify the size of buffers; values of 0 choose suitable defaults.

This routine exists for backward compatibility only. `svc_create()`, `svc_tli_create()`, or `svc_vc_create()` should be used instead. See `rpc_svc_create(3NSL)`.

`svccudp_bufcreate()`

This routine creates a UDP/IP-based RPC service transport, to which it returns a pointer. The transport is associated with the file descriptor `fd`. If `fd` is `RPC_ANYSOCK` then a new file descriptor is created. If the file descriptor is not bound to a local UDP port, then this routine binds it to an arbitrary port. Upon completion, `xprt->xp_fd` is the transport's file descriptor, and `xprt->xp_port` is the transport's port number. Refer to the `File Descriptor` section for more information. This routine returns NULL if it fails.

The user specifies the maximum packet size for sending and receiving UDP-based RPC messages by using the `sendsz` and `recvsz` parameters.

This routine exists for backward compatibility only. `svc_tli_create()`, or `svc_dg_create()` should be used instead. See `rpc_svc_create(3NSL)`.

## rpc\_soc(3NSL)

### svcdp\_create()

This routine creates a UDP/IP-based RPC service transport, to which it returns a pointer. The transport is associated with the file descriptor *fd*, which may be `RPC_ANYSOCK`, in which case a new file descriptor is created. If the file descriptor is not bound to a local UDP port, then this routine binds it to an arbitrary port. Upon completion, *xprt->xp\_fd* is the transport's file descriptor, and *xprt->xp\_port* is the transport's port number. This routine returns NULL if it fails.

Since UDP-based RPC messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take large arguments or return huge results.

This routine exists for backward compatibility only. `svc_create()`, `svc_tli_create()`, or `svc_dg_create()` should be used instead. See `rpc_svc_create(3NSL)`.

### registerrpc()

Register program *prognum*, procedure *procname*, and version *versnum* with the RPC service package. If a request arrives for program *prognum*, version *versnum*, and procedure *procnum*, *procname* is called with a pointer to its parameter(s). *procname* should return a pointer to its static result(s). *inproc* is used to decode the parameters while *outproc* is used to encode the results. This routine returns 0 if the registration succeeded, -1 otherwise.

`svc_run()` must be called after all the services are registered.

This routine exists for backward compatibility only, and it is made obsolete by `rpc_reg()`.

### svc\_register()

Associates *prognum* and *versnum* with the service dispatch procedure, *dispatch*. If *protocol* is 0, the service is not registered with the portmap service. If *protocol* is non-zero, then a mapping of the triple [*prognum*, *versnum*, *protocol*] to *xprt->xp\_port* is established with the local portmap service (generally *protocol* is 0, `IPPROTO_UDP` or `IPPROTO_TCP`). The procedure *dispatch* has the following form:

```
dispatch(struct svc_req *request, SVCXPRT *xprt);
```

The `svc_register()` routine returns one if it succeeds, and 0 otherwise.

This routine exists for backward compatibility only. Enhanced functionality is provided by `svc_reg()`.

### svc\_unregister()

Remove all mapping of the double [*prognum*, *versnum*] to dispatch routines, and of the triple [*prognum*, *versnum*, *all-protocols*] to port number from portmap.

This routine exists for backward compatibility. Enhanced functionality is provided by `svc_unreg()`.

`xdr_authunix_parms()`

Used for describing UNIX credentials. This routine is useful for users who wish to generate these credentials without using the RPC authentication package.

This routine exists for backward compatibility only, and is made obsolete by `xdr_authsys_parms()`. See `rpc_xdr(3NSL)`.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO** `keyserv(1M)`, `rpcbind(1M)`, `rpcinfo(1M)`, `netdir(3NSL)`, `netdir_getbyname(3NSL)`, `rpc(3NSL)`, `rpc_clnt_auth(3NSL)`, `rpc_clnt_calls(3NSL)`, `rpc_clnt_create(3NSL)`, `rpc_svc_calls(3NSL)`, `rpc_svc_create(3NSL)`, `rpc_svc_err(3NSL)`, `rpc_svc_reg(3NSL)`, `rpc_xdr(3NSL)`, `rpcbind(3NSL)`, `secure_rpc(3NSL)`, `select(3C)`, `xdr_authsys_parms(3NSL)`, `libnsl(3LIB)`, `librpcsoc(3LIBUCB)`, `attributes(5)`

**NOTES** These interfaces are unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.

## rpc\_svc\_calls(3NSL)

<b>NAME</b>	rpc_svc_calls, svc_dg_enablecache, svc_done, svc_exit, svc_fdset, svc_freeargs, svc_getargs, svc_getreq_common, svc_getreq_poll, svc_getreqset, svc_getrpcaller, svc_max_pollfd, svc_pollfd, svc_run, svc_sendreply – library routines for RPC servers
<b>DESCRIPTION</b>	<p>These routines are part of the RPC library which allows C language programs to make procedure calls on other machines across the network.</p> <p>These routines are associated with the server side of the RPC mechanism. Some of them are called by the server side dispatch function, while others (such as <code>svc_run()</code>) are called when the server is initiated.</p> <p>In the current implementation, the service transport handle <code>SVCXPRT</code> contains a single data area for decoding arguments and encoding results. Therefore, this structure cannot be freely shared between threads that call functions that do this. However, when a server is operating in the Automatic or User MT modes, a copy of this structure is passed to the service dispatch procedure in order to enable concurrent request processing. Under these circumstances, some routines which would otherwise be unsafe, become safe. These are marked as such. Also marked are routines that are unsafe for MT applications, and are not to be used by such applications.</p>
<b>Routines</b>	<pre>#include &lt;rpc/rpc.h&gt;  int svc_dg_enablecache(SVCXPRT *xpvt, const uint_t cache_size);     This function allocates a duplicate request cache for the service endpoint <i>xpvt</i>, large enough to hold <i>cache_size</i> entries. Once enabled, there is no way to disable caching. This routine returns 1 if space necessary for a cache of the given size was successfully allocated, and 0 otherwise.      This function is safe in MT applications.  int svc_done(SVCXPRT *xpvt);     This function frees resources allocated to service a client request directed to the service endpoint <i>xpvt</i>. This call pertains only to servers executing in the User MT mode. In the User MT mode, service procedures must invoke this call before returning, either after a client request has been serviced, or after an error or abnormal condition that prevents a reply from being sent. After <code>svc_done()</code> is invoked, the service endpoint <i>xpvt</i> should not be referenced by the service procedure. Server multithreading modes and parameters can be set using the <code>rpc_control()</code> call.      This function is safe in MT applications. It will have no effect if invoked in modes other than the User MT mode.  void svc_exit(void);     This function when called by any of the RPC server procedure or otherwise, destroys all services registered by the server and causes <code>svc_run()</code> to return.      If RPC server activity is to be resumed, services must be reregistered with the RPC library either through one of the <code>rpc_svc_create(3NSL)</code> functions, or using <code>xpvt_register(3NSL)</code>.</pre>



`svc_exit()` has global scope and ends all RPC server activity.

`fd_set` `svc_fdset`;

A global variable reflecting the RPC server's read file descriptor bit mask. This is only of interest if service implementors do not call `svc_run()`, but rather do their own asynchronous event processing. This variable is read-only, and it may change after calls to `svc_getreqset()` or any creation routines. Do not pass its address to `select(3C)`! Instead, pass the address of a copy.

MT applications executing in either the Automatic MT mode or the user MT mode should never read this variable. They should use auxiliary threads to do asynchronous event processing.

`svc_fdset` is limited to 1024 file descriptors and is considered obsolete. Use of `svc_pollfd` is recommended instead.

`pollfd_t` `*svc_pollfd`;

A global variable pointing to an array of `pollfd_t` structures reflecting the RPC server's read file descriptor array. This is only of interest if service implementors do not call `svc_run()` but rather do their own asynchronous event processing. This variable is read-only, and it may change after calls to `svc_getreg_poll()` or any creation routines. Do not pass its address to `poll(2)`! Instead, pass the address of a copy.

By default, `svc_pollfd` is limited to 1024 entries. Use `rpc_control(3NSL)` to remove this limitation.

MT applications executing in either the Automatic MT mode or the user MT mode should never be read this variable. They should use auxiliary threads to do asynchronous event processing.

`int` `svc_max_pollfd`;

A global variable containing the maximum length of the `svc_pollfd` array. This variable is read-only, and it may change after calls to `svc_getreg_poll()` or any creation routines.

`bool_t` `svc_freeargs(const SVCXPRT *xprt, const xdrproc_t inproc, caddr_t in)`;

A function macro that frees any data allocated by the RPC/XDR system when it decoded the arguments to a service procedure using `svc_getargs()`. This routine returns `TRUE` if the results were successfully freed, and `FALSE` otherwise.

This function macro is safe in MT applications utilizing the Automatic or User MT modes.

`bool_t` `svc_getargs(const SVCXPRT *xprt, const xdrproc_t inproc, caddr_t in)`;

A function macro that decodes the arguments of an RPC request associated with the RPC service transport handle `xprt`. The parameter `in` is the address where the arguments will be placed; `inproc` is the XDR routine used to decode the arguments. This routine returns `TRUE` if decoding succeeds, and `FALSE` otherwise.

## rpc\_svc\_calls(3NSL)

This function macro is safe in MT applications utilizing the Automatic or User MT modes.

```
void svc_getreq_common(const int fd);
```

This routine is called to handle a request on the given file descriptor.

```
void svc_getreq_poll(struct pollfd *pfdp, const int pollretval) ;
```

This routine is only of interest if a service implementor does not call `svc_run()`, but instead implements custom asynchronous event processing. It is called when `poll(2)` has determined that an RPC request has arrived on some RPC file descriptors; *pollretval* is the return value from `poll(2)` and *pfdp* is the array of *pollfd* structures on which the `poll(2)` was done. It is assumed to be an array large enough to contain the maximal number of descriptors allowed.

This function macro is unsafe in MT applications.

```
void svc_getreqset(fd_set *rdfds);
```

This routine is only of interest if a service implementor does not call `svc_run()`, but instead implements custom asynchronous event processing. It is called when `select(3C)` has determined that an RPC request has arrived on some RPC file descriptors; *rdfds* is the resultant read file descriptor bit mask. The routine returns when all file descriptors associated with the value of *rdfds* have been serviced.

This function macro is unsafe in MT applications.

```
struct netbuf *svc_getrpccaller(const SVCXPRT *xprt);
```

The approved way of getting the network address of the caller of a procedure associated with the RPC service transport handle *xprt*.

This function macro is safe in MT applications.

```
void svc_run(void);
```

This routine never returns. In single threaded mode, it waits for RPC requests to arrive, and calls the appropriate service procedure using `svc_getreq_poll()` when one arrives. This procedure is usually waiting for the `poll(2)` library call to return.

Applications executing in the Automatic or User MT modes should invoke this function exactly once. In the Automatic MT mode, it will create threads to service client requests. In the User MT mode, it will provide a framework for service developers to create and manage their own threads for servicing client requests.

```
bool_t svc_sendreply(const SVCXPRT *xprt, const xdrproc_t outproc, const caddr_t out);
```

Called by an RPC service's dispatch routine to send the results of a remote procedure call. The parameter *xprt* is the request's associated transport handle; *outproc* is the XDR routine which is used to encode the results; and *out* is the address of the results. This routine returns `TRUE` if it succeeds, `FALSE` otherwise.

This function macro is safe in MT applications utilizing the Automatic or User MT modes.

rpc\_svc\_calls(3NSL)

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	See NOTES below.

**SEE ALSO** rpcgen(1), poll(2), rpc(3NSL), rpc\_control(3NSL), rpc\_svc\_create(3NSL), rpc\_svc\_err(3NSL), rpc\_svc\_reg(3NSL), select(3C), xpirt\_register(3NSL), attributes(5)

**NOTES** svc\_dg\_enablecache() and svc\_getrpccaller() are safe in multithreaded applications. svc\_freeargs(), svc\_getargs(), and svc\_sendreply() are safe in MT applications utilizing the Automatic or User MT modes. svc\_getreq\_common(), svc\_getreqset(), and svc\_getreq\_poll() are unsafe in multithreaded applications and should be called only from the main thread.

## rpc\_svc\_create(3NSL)

<b>NAME</b>	rpc_svc_create, svc_control, svc_create, svc_destroy, svc_dg_create, svc_fd_create, svc_raw_create, svc_tli_create, svc_tp_create, svc_vc_create, svc_door_create – library routines for the creation of server handles
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpc.h&gt;  bool_t <b>svc_control</b>(SVCXPRT *svc, const uint_t req, void *info);  int <b>svc_create</b>(const void (*dispatch) const struct svc_req *, const     SVCXPRT *, const rpcprog_t prognum, const rpcvers_t versnum,     const char *nettype);  void <b>svc_destroy</b>(SVCXPRT *xpirt);  SVCXPRT *<b>svc_dg_create</b>(const int fildes, const uint_t sendsz, const     uint_t recvsz);  SVCXPRT *<b>svc_fd_create</b>(const int fildes, const uint_t sendsz, const     uint_t recvsz);  SVCXPRT *<b>svc_raw_create</b>(void);  SVCXPRT *<b>svc_tli_create</b>(const int fildes, const struct netconfig     *netconf, const struct t_bind *bind_addr, const uint_t sendsz,     const uint_t recvsz);  SVCXPRT *<b>svc_tp_create</b>(const void (*dispatch) const struct svc_req     *, const SVCXPRT *), const rpcprog_t prognum, const rpcvers_t     versnum, const struct netconfig *netconf);  SVCXPRT *<b>svc_vc_create</b>(const int fildes, const uint_t sendsz, const     uint_t recvsz);  SVCXPRT *<b>svc_door_create</b>(void (*dispatch) (struct svc_req *, SVCXPRT     *), const rpcprog_t prognum, const rpcvers_t versnum, const     uint_t sendsz);</pre>
<b>DESCRIPTION</b>	These routines are part of the RPC library which allows C language programs to make procedure calls on servers across the network. These routines deal with the creation of service handles. Once the handle is created, the server can be invoked by calling <code>svc_run()</code> .
<b>Routines</b>	See <code>rpc(3NSL)</code> for the definition of the <code>SVCXPRT</code> data structure.  <code>svc_control()</code> A function to change or retrieve information about a service object. <i>req</i> indicates the type of operation and <i>info</i> is a pointer to the information. The supported values of <i>req</i> , their argument types, and what they do are:  SVCGET_VERSQUIET If a request is received for a program number served by this server but the version number is outside the range registered with the server, an

## rpc\_svc\_create(3NSL)

RPC\_PROGVERSMISMATCH error will normally be returned. *info* should be a pointer to an integer. Upon successful completion of the SVCGET\_VERSQUIET request, *\*info* contains an integer which describes the server's current behavior: 0 indicates normal server behavior, that is, an RPC\_PROGVERSMISMATCH error will be returned. 1 indicates that the out of range request will be silently ignored.

### SVCSET\_VERSQUIET

If a request is received for a program number served by this server but the version number is outside the range registered with the server, an RPC\_PROGVERSMISMATCH error will normally be returned. It is sometimes desirable to change this behavior. *info* should be a pointer to an integer which is either 0, indicating normal server behavior and an RPC\_PROGVERSMISMATCH error will be returned, or 1, indicating that the out of range request should be silently ignored.

### SVCGET\_XID

Returns the transaction ID of connection-oriented and connectionless transport service calls. The transaction ID assists in uniquely identifying client requests for a given RPC version, program number, procedure, and client. The transaction ID is extracted from the service transport handle *svc*. *info* must be a pointer to an unsigned long. Upon successful completion of the SVCGET\_XID request, *\*info* contains the transaction ID. Note that rendezvous and raw service handles do not define a transaction ID. Thus, if the service handle is of rendezvous or raw type, and the request is of type SVCGET\_XID, *svc\_control()* will return FALSE. Note also that the transaction ID read by the server can be set by the client through the suboption CLSET\_XID in *clnt\_control()*. See *clnt\_create(3NSL)*

### SVCSET\_RECVERRHANDLER

Attaches or detaches a disconnection handler to the service handle, *svc*, that will be called when a transport error arrives during the reception of a request or when the server is waiting for a request and the connection shuts down. This handler is only useful for a connection oriented service handle.

## rpc\_svc\_create(3NSL)

	<p><i>*info</i> contains the address of the error handler to attach, or NULL to detach a previously defined one. The error handler has two arguments. It has a pointer to the erroneous service handle. It also has an integer that indicates if the full service is closed (when equal to zero), or that only one connection on this service is closed (when not equal to zero).</p> <pre>void handler (const SVCXPRT *svc, const bool_t isAConnection);</pre> <p>With the service handle address, <i>svc</i>, the error handler is able to detect which connection has failed and to begin an error recovery process. The error handler can be called by multiple threads and should be implemented in an MT-safe way.</p> <p>SVCGET_RECVERRHANDLER Upon successful completion of the SVCGET_RECVERRHANDLER request, <i>*info</i> contains the address of the handler for receiving errors. Upon failure, <i>*info</i> contains NULL.</p> <p>This routine returns TRUE if the operation was successful. Otherwise, it returns false.</p>
<code>svc_create()</code>	<p><code>svc_create()</code> creates server handles for all the transports belonging to the class <i>nettype</i>.</p> <p><i>nettype</i> defines a class of transports which can be used for a particular application. The transports are tried in left to right order in NETPATH variable or in top to bottom order in the netconfig database. If <i>nettype</i> is NULL, it defaults to netpath.</p> <p><code>svc_create()</code> registers itself with the rpcbind service (see rpcbind(1M)). <i>dispatch</i> is called when there is a remote procedure call for the given <i>prognum</i> and <i>versnum</i>; this requires calling <code>svc_run()</code> (see <code>svc_run()</code> in <code>rpc_svc_reg(3NSL)</code>). If <code>svc_create()</code> succeeds, it returns the number of server handles it created, otherwise it returns 0 and an error message is logged.</p>
<code>svc_destroy()</code>	<p>A function macro that destroys the RPC service handle <i>xprt</i>. Destruction usually involves deallocation of private data structures, including <i>xprt</i> itself. Use of <i>xprt</i> is undefined after calling this routine.</p>
<code>svc_dg_create()</code>	<p>This routine creates a connectionless RPC service handle, and returns a pointer to it. This routine returns NULL if it fails, and an error message is logged. <i>sendsz</i></p>

rpc\_svc\_create(3NSL)

and *recvsz* are parameters used to specify the size of the buffers. If they are 0, suitable defaults are chosen. The file descriptor *fildev* should be open and bound. The server is not registered with `rpcbind(1M)`.

Warning: since connectionless-based RPC messages can only hold limited amount of encoded data, this transport cannot be used for procedures that take large arguments or return huge results.

`svc_fd_create()`

This routine creates a service on top of an open and bound file descriptor, and returns the handle to it. Typically, this descriptor is a connected file descriptor for a connection-oriented transport. *sendsz* and *recvsz* indicate sizes for the send and receive buffers. If they are 0, reasonable defaults are chosen. This routine returns NULL if it fails, and an error message is logged.

`svc_raw_create()`

This routine creates an RPC service handle and returns a pointer to it. The transport is really a buffer within the process's address space, so the corresponding RPC client should live in the same address space; (see `clnt_raw_create()` in `rpc_clnt_create(3NSL)`). This routine allows simulation of RPC and acquisition of RPC overheads (such as round trip times), without any kernel and networking interference. This routine returns NULL if it fails, and an error message is logged.

Note: `svc_run()` should not be called when the raw interface is being used.

`svc_tli_create()`

This routine creates an RPC server handle, and returns a pointer to it. *fildev* is the file descriptor on which the service is listening. If *fildev* is `RPC_ANYFD`, it opens a file descriptor on the transport specified by *netconf*. If the file descriptor is unbound and *bindaddr* is non-null *fildev* is bound to the address specified by *bindaddr*, otherwise *fildev* is bound to a default address chosen by the transport. In the case where the default address is chosen, the number of outstanding connect requests is set to 8 for connection-oriented transports. The user may specify the size of the send and receive buffers with the parameters *sendsz* and *recvsz*; values of 0 choose suitable defaults. This routine returns NULL if it fails, and an error message is logged. The server is not registered with the `rpcbind(1M)` service.

`svc_tp_create()`

`svc_tp_create()` creates a server handle for the network specified by *netconf*, and registers itself with

## rpc\_svc\_create(3NSL)

the `rpcbind` service. `dispatch` is called when there is a remote procedure call for the given `prognum` and `versnum`; this requires calling `svc_run()`.

`svc_tp_create()` returns the service handle if it succeeds, otherwise a `NULL` is returned and an error message is logged.

### `svc_vc_create()`

This routine creates a connection-oriented RPC service and returns a pointer to it. This routine returns `NULL` if it fails, and an error message is logged. The users may specify the size of the send and receive buffers with the parameters `sendsz` and `recvsz`; values of 0 choose suitable defaults. The file descriptor `fildev` should be open and bound. The server is not registered with the `rpcbind(1M)` service.

### `svc_door_create()`

This routine creates an RPC server handle over doors and returns a pointer to it. Doors is a transport mechanism that facilitates fast data transfer between processes on the same machine. for the given program The user may set the size of the send buffer with the parameter `sendsz`. If `sendsz` is 0, the corresponding default buffer size is 16 Kbyte. If successful, the `svc_door_create()` routine returns the service handle. Otherwise it returns `NULL` and sets a value for `rpc_createerr`. The server is not registered with `rpcbind(1M)`.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	All
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Evolving
MT-Level	MT-Safe

**SEE ALSO** `rpcbind(1M)`, `rpc(3NSL)`, `rpc_clnt_create(3NSL)`, `rpc_svc_calls(3NSL)`, `rpc_svc_err(3NSL)`, `rpc_svc_reg(3NSL)`, `attributes(5)`



<b>NAME</b>	rpc_svc_err, svcerr_auth, svcerr_decode, svcerr_noproc, svcerr_noprog, svcerr_progvers, svcerr_systemerr, svcerr_weakauth – library routines for server side remote procedure call errors
<b>DESCRIPTION</b>	<p>These routines are part of the RPC library which allows C language programs to make procedure calls on other machines across the network.</p> <p>These routines can be called by the server side dispatch function if there is any error in the transaction with the client.</p>
<b>Routines</b>	<p>See <code>rpc(3NSL)</code> for the definition of the <code>SVCXPRT</code> data structure.</p> <pre>#include &lt;rpc/rpc.h&gt;</pre> <p><code>void svcerr_auth(const SVCXPRT *xpvt, const enum auth_stat why);</code>  Called by a service dispatch routine that refuses to perform a remote procedure call due to an authentication error.</p> <p><code>void svcerr_decode(const SVCXPRT *xpvt);</code>  Called by a service dispatch routine that cannot successfully decode the remote parameters (see <code>svc_getargs()</code> in <code>rpc_svc_reg(3NSL)</code>).</p> <p><code>void svcerr_noproc(const SVCXPRT *xpvt);</code>  Called by a service dispatch routine that does not implement the procedure number that the caller requests.</p> <p><code>void svcerr_noprog(const SVCXPRT *xpvt);</code>  Called when the desired program is not registered with the RPC package. Service implementors usually do not need this routine.</p> <p><code>void svcerr_progvers(const SVCXPRT *xpvt, const rpcvers_t low_vers, const rpcvers_t high_vers);</code>  Called when the desired version of a program is not registered with the RPC package. <i>low_vers</i> is the lowest version number, and <i>high_vers</i> is the highest version number. Service implementors usually do not need this routine.</p> <p><code>void svcerr_systemerr(const SVCXPRT *xpvt);</code>  Called by a service dispatch routine when it detects a system error not covered by any particular protocol. For example, if a service can no longer allocate storage, it may call this routine.</p> <p><code>void svcerr_weakauth(const SVCXPRT *xpvt);</code>  Called by a service dispatch routine that refuses to perform a remote procedure call due to insufficient (but correct) authentication parameters. The routine calls <code>svcerr_auth(xpvt, AUTH_TOOWEAK)</code>.</p>
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

rpc\_svc\_err(3NSL)

**SEE ALSO** | `rpc(3NSL)`, `rpc_svc_calls(3NSL)`, `rpc_svc_create(3NSL)`,  
`rpc_svc_reg(3NSL)`, `attributes(5)`

<b>NAME</b>	rpc_svc_input, svc_add_input, svc_remove_input – declare or remove a callback on a file descriptor						
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/rpc.h&gt;  typedef void (*svc_callback_t) (svc_input_id_t id, int fd, unsigned     int events, void *cookie);  svc_input_id_t svc_add_input(int fd, unsigned int revents,     svc_callback_t callback, void *cookie);  int svc_remove_input(svc_input_t id);</pre>						
<b>DESCRIPTION</b>	The following RPC routines are used to declare or remove a callback on a file descriptor.						
<b>Routines</b>	<p>See rpc(3NSL) for the definition of the SVCXPRT data structure.</p> <p><b>svc_add_input ()</b>  This function is used to register a <i>callback</i> function on a file descriptor, <i>fd</i>. The file descriptor, <i>fd</i>, is the first parameter to be passed to <i>svc_add_input ()</i>. This <i>callback</i> function will be automatically called if any of the events specified in the <i>events</i> parameter occur on this descriptor. The <i>events</i> parameter is used to specify when the callback is invoked. This parameter is a mask of poll events to which the user wants to listen. See <i>poll(2)</i> for further details of the events that can be specified.</p> <p>The callback to be invoked is specified using the <i>callback</i> parameter. The <i>cookie</i> parameter can be used to pass any data to the <i>callback</i> function. This parameter is a user-defined value which is passed as an argument to the <i>callback</i> function, and it is not used by the Sun RPC library itself.</p> <p>Several callbacks can be registered on the same file descriptor as long as each callback registration specifies a separate set of event flags.</p> <p>The <i>callback</i> function is called with the registration <i>id</i>, the <i>fd</i> file descriptor, an <i>revents</i> value, which is a bitmask of all events concerning the file descriptor, and the <i>cookie</i> user-defined value.</p> <p>Upon successful completion, the function returns a unique identifier for this registration, that can be used later to remove this callback. Upon failure, -1 is returned and <i>errno</i> is set to indicate the error.</p> <p>The <i>svc_add_input ()</i> function will fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The <i>fd</i> or <i>events</i> parameters are invalid.</td> </tr> <tr> <td>EEXIST</td> <td>A callback is already registered to the file descriptor with one of the specified events.</td> </tr> <tr> <td>ENOMEM</td> <td>Memory is exhausted.</td> </tr> </table>	EINVAL	The <i>fd</i> or <i>events</i> parameters are invalid.	EEXIST	A callback is already registered to the file descriptor with one of the specified events.	ENOMEM	Memory is exhausted.
EINVAL	The <i>fd</i> or <i>events</i> parameters are invalid.						
EEXIST	A callback is already registered to the file descriptor with one of the specified events.						
ENOMEM	Memory is exhausted.						

## rpc\_svc\_input(3NSL)

`svc_remove_input()`

This function is used to unregister a callback function on a file descriptor, *fd*. The *id* parameter specifies the registration to be removed.

Upon successful completion, the function returns zero. Upon failure, -1 is returned and `errno` is set to indicate the error.

The `svc_remove_input()` function will fail if:

`EINVAL` The *id* parameter is invalid.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	All
Availability	SUNWcsl (32-bit)
	SUNWcslx (64-bit)
Interface Stability	Evolving
MT-Level	MT-Safe

**SEE ALSO** `poll(2)`, `rpc(3NSL)`, `attributes(5)`

<b>NAME</b>	rpc_svc_reg, rpc_reg, svc_reg, svc_unreg, svc_auth_reg, xprt_register, xprt_unregister – library routines for registering servers
<b>DESCRIPTION</b>	These routines are a part of the RPC library which allows the RPC servers to register themselves with <code>rpcbind()</code> (see <code>rpcbind(1M)</code> ), and associate the given program and version number with the dispatch function. When the RPC server receives a RPC request, the library invokes the dispatch routine with the appropriate arguments.
<b>Routines</b>	<p>See <code>rpc(3NSL)</code> for the definition of the <code>SVCXPRT</code> data structure.</p> <pre>#include &lt;rpc/rpc.h&gt;</pre> <p><code>bool_t rpc_reg(const rpcprog_t prognum, const rpcvers_t versnum, const rpcproc_t procnum, char * (*procname)(), const xdrproc_t inproc, const xdrproc_t outproc, const char *nettype);</code>  Register program <i>prognum</i>, procedure <i>procname</i>, and version <i>versnum</i> with the RPC service package. If a request arrives for program <i>prognum</i>, version <i>versnum</i>, and procedure <i>procnum</i>, <i>procname</i> is called with a pointer to its parameter(s); <i>procname</i> should return a pointer to its <code>static</code> result(s). The <i>arg</i> parameter to <i>procname</i> is a pointer to the (decoded) procedure argument. <i>inproc</i> is the XDR function used to decode the parameters while <i>outproc</i> is the XDR function used to encode the results. Procedures are registered on all available transports of the class <i>nettype</i>. See <code>rpc(3NSL)</code>. This routine returns 0 if the registration succeeded, -1 otherwise.</p> <p><code>int svc_reg(const SVCXPRT *xprt, const rpcprog_t prognum, const rpcvers_t versnum, const void (*dispatch)(), const struct netconfig *netconf);</code>  Associates <i>prognum</i> and <i>versnum</i> with the service dispatch procedure, <i>dispatch</i>. If <i>netconf</i> is <code>NULL</code>, the service is not registered with the <code>rpcbind</code> service. For example, if a service has already been registered using some other means, such as <code>inetd</code> (see <code>inetd(1M)</code>), it will not need to be registered again. If <i>netconf</i> is non-zero, then a mapping of the triple [<i>prognum</i>, <i>versnum</i>, <i>netconf</i>-&gt;] to <i>xprt</i>-&gt; <i>xp_ltaddr</i> is established with the local <code>rpcbind</code> service.</p> <p>The <code>svc_reg()</code> routine returns 1 if it succeeds, and 0 otherwise.</p> <p><code>void svc_unreg(const rpcprog_t prognum, const rpcvers_t versnum);</code>  Remove from the <code>rpcbind</code> service, all mappings of the triple [<i>prognum</i>, <i>versnum</i>, <i>all-transports</i>] to network address and all mappings within the RPC service package of the double [<i>prognum</i>, <i>versnum</i>] to dispatch routines.</p> <p><code>int svc_auth_reg(const int cred_flavor, const enum auth_stat (*handler)());</code>  Registers the service authentication routine <i>handler</i> with the dispatch mechanism so that it can be invoked to authenticate RPC requests received with authentication type <i>cred_flavor</i>. This interface allows developers to add new authentication types to their RPC applications without needing to modify the libraries. Service implementors usually do not need this routine.</p>

## rpc\_svc\_reg(3NSL)

Typical service application would call `svc_auth_reg()` after registering the service and prior to calling `svc_run()`. When needed to process an RPC credential of type `cred_flavor`, the *handler* procedure will be called with two parameters (`struct svc_req *rqst`, `struct rpc_msg *msg`) and is expected to return a valid enum `auth_stat` value. There is no provision to change or delete an authentication handler once registered.

The `svc_auth_reg()` routine returns 0 if the registration is successful, 1 if `cred_flavor` already has an authentication handler registered for it, and -1 otherwise.

`void xprt_register(const SVCXPRT *xprt);`

After RPC service transport handle `xprt` is created, it is registered with the RPC service package. This routine modifies the global variable `svc_fdset` (see `rpc_svc_calls(3NSL)`). Service implementors usually do not need this routine.

`void xprt_unregister(const SVCXPRT *xprt);`

Before an RPC service transport handle `xprt` is destroyed, it unregisters itself with the RPC service package. This routine modifies the global variable `svc_fdset` (see `rpc_svc_calls(3NSL)`). Service implementors usually do not need this routine.

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

### SEE ALSO

`inetd(1M)`, `rpcbind(1M)`, `rpc(3NSL)`, `rpc_svc_calls(3NSL)`, `rpc_svc_create(3NSL)`, `rpc_svc_err(3NSL)`, `rpcbind(3NSL)`, `select(3C)`, `attributes(5)`

<b>NAME</b>	rpc_xdr, xdr_accepted_reply, xdr_authsys_parms, xdr_callhdr, xdr_callmsg, xdr_opaque_auth, xdr_rejected_reply, xdr_replymsg – XDR library routines for remote procedure calls
<b>SYNOPSIS</b>	<pre> bool_t xdr_accepted_reply(XDR *xdrs, const struct accepted_reply     *ar) ;  bool_t xdr_authsys_parms(XDR *xdrs, struct authsys_parms *aupp) ; void xdr_callhdr(XDR *xdrs, struct rpc_msg *chdr) ; bool_t xdr_callmsg(XDR *xdrs, struct rpc_msg *cmsg) ; bool_t xdr_opaque_auth(XDR *xdrs, struct opaque_auth *ap) ; bool_t xdr_rejected_reply(XDR *xdrs, const struct rejected_reply     *rr) ;  bool_t xdr_replymsg(XDR *xdrs, const struct rpc_msg *rmsg) ; </pre>
<b>DESCRIPTION</b>	These routines are used for describing the RPC messages in XDR language. They should normally be used by those who do not want to use the RPC package directly. These routines return TRUE if they succeed, FALSE otherwise.
<b>Routines</b>	<p>See rpc(3NSL) for the definition of the XDR data structure.</p> <pre> #include &lt;rpc/rpc.h&gt; </pre> <p><b>xdr_accepted_reply()</b> Used to translate between RPC reply messages and their external representation. It includes the status of the RPC call in the XDR language format. In the case of success, it also includes the call results.</p> <p><b>xdr_authsys_parms()</b> Used for describing UNIX operating system credentials. It includes machine-name, uid, gid list, etc.</p> <p><b>xdr_callhdr()</b> Used for describing RPC call header messages. It encodes the static part of the call message header in the XDR language format. It includes information such as transaction ID, RPC version number, program and version number.</p> <p><b>xdr_callmsg()</b> Used for describing RPC call messages. This includes all the RPC call information such as transaction ID, RPC version number, program number, version number, authentication information, etc. This is normally used by servers to determine information about the client RPC call.</p> <p><b>xdr_opaque_auth()</b> Used for describing RPC opaque authentication information messages.</p>

rpc\_xdr(3NSL)

`xdr_rejected_reply()`

Used for describing RPC reply messages. It encodes the rejected RPC message in the XDR language format. The message could be rejected either because of version number mis-match or because of authentication errors.

`xdr_replymsg()`

Used for describing RPC reply messages. It translates between the RPC reply message and its external representation. This reply could be either an acceptance, rejection or NULL.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `rpc(3NSL)`, `xdr(3NSL)`, `attributes(5)`



<b>NAME</b>	rstat, havedisk – get performance data from remote kernel				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lrpcsvc [ library ... ] #include &lt;rpc/rpc.h&gt; #include &lt;rpcsvc/rstat.h&gt;  enum clnt_stat <b>rstat</b>(char *host, struct statstime *statp); int <b>havedisk</b>(char *host);</pre>				
<b>PROTOCOL</b>	/usr/include/rpcsvc/rstat.x				
<b>DESCRIPTION</b>	<p>These routines require that the <code>rpc.rstatd(1M)</code> daemon be configured and available on the remote system indicated by <i>host</i>. The <code>rstat()</code> protocol is used to gather statistics from remote kernel. Statistics will be available on items such as paging, swapping, and cpu utilization.</p> <p><code>rstat()</code> fills in the <code>statstime</code> structure <i>statp</i> for <i>host</i>. <i>statp</i> must point to an allocated <code>statstime</code> structure. <code>rstat()</code> returns <code>RPC_SUCCESS</code> if it was successful; otherwise a <code>enum clnt_stat</code> is returned which can be displayed using <code>clnt_perrno(3NSL)</code>.</p> <p><code>havedisk()</code> returns 1 if <i>host</i> has disk, 0 if it does not, and -1 if this cannot be determined.</p> <p>The following XDR routines are available in <code>librpcsvc</code>:</p> <pre>xdr_statstime xdr_statsvar</pre>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>rup(1)</code> , <code>rpc.rstatd(1M)</code> , <code>rpc_clnt_calls(3NSL)</code> , <code>attributes(5)</code>				

## rusers(3RPC)

<b>NAME</b>	rusers, rnusers – return information about users on remote machines				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lrpcsvc [ library ... ] #include &lt;rpc/rpc.h&gt; #include &lt;rpcsvc/rusers.h&gt;  enum clnt_stat <b>rusers</b>(char *host, struct utmpidlearr *up); int <b>rnusers</b>(char *host);</pre>				
<b>PROTOCOL</b>	/usr/include/rpcsvc/rusers.x				
<b>DESCRIPTION</b>	<p>These routines require that the <code>rpc.rusersd(1M)</code> daemon be configured and available on the remote system indicated by <i>host</i>. The <code>rusers()</code> protocol is used to retrieve information about users logged in on the remote system.</p> <p><code>rusers()</code> fills the <code>utmpidlearr</code> structure with data about <i>host</i>, and returns 0 if successful. <i>up</i> must point to an allocated <code>utmpidlearr</code> structure. If <code>rusers()</code> returns successful it will have allocated data structures within the <i>up</i> structure, which should be freed with <code>xdr_free(3NSL)</code> when you no longer need them:</p> <pre>xdr_free(xdr_utmpidlearr, up);</pre> <p>On error, the returned value can be interpreted as an <code>enum clnt_stat</code> and can be displayed with <code>clnt_perror(3NSL)</code> or <code>clnt_sperrno(3NSL)</code>.</p> <p>See the header <code>&lt;rpcsvc/rusers.h&gt;</code> for a definition of <code>struct utmpidlearr</code>.</p> <p><code>rnusers()</code> returns the number of users logged on to <i>host</i> (-1 if it cannot determine that number).</p> <p>The following XDR routines are available in <code>librpcsvc</code>:</p> <pre>xdr_utmpidlearr</pre>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>rusers(1)</code> , <code>rpc.rusersd(1M)</code> , <code>rpc_clnt_calls(3NSL)</code> , <code>xdr_free(3NSL)</code> , <code>attributes(5)</code>				

<b>NAME</b>	rwall – write to specified remote machines				
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lrpcsvc [ library ... ] #include &lt;rpc/rpc.h&gt; #include &lt;rpcsvc/rwall.h&gt;  enum clnt_stat <b>rwall</b>(char *host, char *msg);</pre>				
<b>PROTOCOL</b>	/usr/include/rpcsvc/rwall.x				
<b>DESCRIPTION</b>	<p>These routines require that the <code>rpc.rwalld(1M)</code> daemon be configured and available on the remote system indicated by <i>host</i>.</p> <p><code>rwall()</code> executes <code>wall(1M)</code> on <i>host</i>. The <code>rpc.rwalld</code> process on <i>host</i> prints <i>msg</i> to all users logged on to that system. <code>rwall()</code> returns <code>RPC_SUCCESS</code> if it was successful; otherwise a <code>enum clnt_stat</code> is returned which can be displayed using <code>clnt_perrno(3NSL)</code>.</p>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	<code>rpc.rwalld(1M)</code> , <code>wall(1M)</code> , <code>rpc_clnt_calls(3NSL)</code> , <code>attributes(5)</code>				

## secure\_rpc(3NSL)

<b>NAME</b>	secure_rpc, authdes_getucrd, authdes_seccreate, getnetname, host2netname, key_decryptsession, key_encryptsession, key_gendes, key_setsecret, key_secretkey_is_set, netname2host, netname2user, user2netname – library routines for secure remote procedure calls						
<b>SYNOPSIS</b>	<pre>int authdes_getucrd(const struct authdes_cred *adc, uid_t *uidp, gid_t *gidp, short *gidlenp, gid_t *gidlist);  AUTH *authdes_seccreate(const char *name, const uint_t window, const char *timehost, const des_block *ckey);  int getnetname(char name [MAXNETNAMELEN+1]);  int host2netname(char name [MAXNETNAMELEN+1], const char *host, const char *domain);  int key_decryptsession(const char *remotename, des_block *deskey); int key_encryptsession(const char *remotename, des_block *deskey); int key_gendes(des_block *deskey); int key_setsecret(const char *key); int key_secretkey_is_set(void); int netname2host(const char *name, char *host, const int hostlen); int netname2user(const char *name, uid_t *uidp, gid_t *gidp, int *gidlenp, gid_t *gidlist [NGRPS]); int user2netname(char name [MAXNETNAMELEN+1], const uid_t uid, const char *domain);</pre>						
<b>DESCRIPTION</b>	<p>RPC library routines allow C programs to make procedure calls on other machines across the network.</p> <p>RPC supports various authentication flavors. Among them are:</p> <table><tr><td>AUTH_NONE</td><td>No authentication (none).</td></tr><tr><td>AUTH_SYS</td><td>Traditional UNIX-style authentication.</td></tr><tr><td>AUTH_DES</td><td>DES encryption-based authentication.</td></tr></table> <p>The <code>authdes_getucrd()</code> and <code>authdes_seccreate()</code> routines implement the AUTH_DES authentication flavor. The keyserver daemon <code>keyserv</code> (see <code>keyserv(1M)</code>) must be running for the AUTH_DES authentication system to work, and <code>keylogin(1)</code> must have been run. Only the AUTH_DES style of authentication is discussed here. For information about the AUTH_NONE and AUTH_SYS styles of authentication, refer to <code>rpc_clnt_auth(3NSL)</code>.</p> <p>The routines documented on this page are MT-Safe. See the man pages for the other authentication styles for their MT-level.</p>	AUTH_NONE	No authentication (none).	AUTH_SYS	Traditional UNIX-style authentication.	AUTH_DES	DES encryption-based authentication.
AUTH_NONE	No authentication (none).						
AUTH_SYS	Traditional UNIX-style authentication.						
AUTH_DES	DES encryption-based authentication.						
<b>Routines</b>	See <code>rpc(3NSL)</code> for the definition of the AUTH data structure.						

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

#### authdes\_getucred()

`authdes_getucred()` is the first of the two routines which interface to the RPC secure authentication system known as AUTH\_DES. The second is `authdes_seccreate()`, below. `authdes_getucred()` is used on the server side for converting an AUTH\_DES credential, which is operating system independent, into an AUTH\_SYS credential. This routine returns 1 if it succeeds, 0 if it fails.

*\*uidp* is set to the user's numerical ID associated with *adc*. *\*gidp* is set to the numerical ID of the user's group. *\*gidlist* contains the numerical IDs of the other groups to which the user belongs. *\*gidlenp* is set to the number of valid group ID entries in *\*gidlist* (see `netname2user()`, below).

Warning: `authdes_getucred()` will fail if the `authdes_cred` structure was created with the netname of a host. In such a case, `netname2host()` should be used on the host netname in the `authdes_cred` structure to get the host name.

#### authdes\_seccreate()

`authdes_seccreate()`, the second of two AUTH\_DES authentication routines, is used on the client side to return an authentication handle that will enable the use of the secure authentication system. The first parameter *name* is the network name, or *netname*, of the owner of the server process. This field usually represents a hostname derived from the utility routine `host2netname()`, but could also represent a user name using `user2netname()`, described below.

The second field is *window* on the validity of the client credential, given in seconds. If the difference in time between the client's clock and the server's clock exceeds *window*, the server will reject the client's credentials, and the clock will have to be resynchronized. A small window is more secure than a large one, but choosing too small of a window will increase the frequency of resynchronizations because of clock drift.

The third parameter, *timehost*, the host's name, is optional. If it is NULL, then the authentication system will assume that the local clock is always in sync with the *timehost* clock, and will not attempt resynchronizations. If a *timehost* is supplied, however, then the system will consult with the remote time service whenever resynchronization is required. This parameter is usually the name of the host on which the server is running.

The final parameter *ckey* is also optional. If it is NULL, then the authentication system will generate a random DES key to be used for the encryption of credentials. If *ckey* is supplied, then it will be used instead.

If `authdes_seccreate()` fails, it returns NULL.

## secure\_rpc(3NSL)

### getnetname ()

getnetname () returns the unique, operating system independent netname of the caller in the fixed-length array *name*. Returns 1 if it succeeds, and 0 if it fails.

### host2netname ()

Convert from a domain-specific hostname *host* to an operating system independent netname. Returns 1 if it succeeds, and 0 if it fails. Inverse of netname2host (). If *domain* is NULL, host2netname () uses the default domain name of the machine. If *host* is NULL, it defaults to that machine itself. If *domain* is NULL and *host* is a NIS name like "host1.ssi.sun.com," host2netname () uses the domain "ssi.sun.com" rather than the default domain name of the machine.

### key\_decryptsession ()

key\_decryptsession () is an interface to the keyserver daemon, which is associated with RPC's secure authentication system (AUTH\_DES authentication).

User programs rarely need to call it, or its associated routines

key\_encryptsession (), key\_gendes (), and key\_setsecret ().

key\_decryptsession () takes a server netname *remotename* and a DES key *deskey*, and decrypts the key by using the the public key of the the server and the secret key associated with the effective UID of the calling process. It is the inverse of key\_encryptsession ().

### key\_encryptsession ()

key\_encryptsession () is a keyserver interface routine. It takes a server netname *remotename* and a DES key *deskey*, and encrypts it using the public key of the the server and the secret key associated with the effective UID of the calling process. It is the inverse of key\_decryptsession (). This routine returns 0 if it succeeds, -1 if it fails.

### key\_gendes ()

key\_gendes () is a keyserver interface routine. It is used to ask the keyserver for a secure conversation key. Choosing one at random is usually not good enough, because the common ways of choosing random numbers, such as using the current time, are very easy to guess. This routine returns 0 if it succeeds, -1 if it fails.

### key\_setsecret ()

key\_setsecret () is a keyserver interface routine. It is used to set the key for the effective UID of the calling process. This routine returns 0 if it succeeds, -1 if it fails.

### key\_secretkey\_is\_set ()

key\_secretkey\_is\_set () is a keyserver interface routine that may be used to determine whether a key has been set for the effective UID of the calling process. If the keyserver has a key stored for the effective UID of the calling process, this routine returns 1. Otherwise it returns 0.

### netname2host ()

Convert from an operating system independent netname *name* to a domain-specific hostname *host*. *hostlen* is the maximum size of *host*. Returns 1 if it succeeds, and 0 if it fails. Inverse of host2netname ().

netname2user ()

Convert from an operating system independent netname to a domain-specific user ID. Returns 1 if it succeeds, and 0 if it fails. Inverse of user2netname ().

\*uidp is set to the user's numerical ID associated with *name*. \*gidp is set to the numerical ID of the user's group. *gidlist* contains the numerical IDs of the other groups to which the user belongs. \*gidlenp is set to the number of valid group ID entries in *gidlist*.

user2netname ()

Convert from a domain-specific username to an operating system independent netname. Returns 1 if it succeeds, and 0 if it fails. Inverse of netname2user ().

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** chkey(1), keylogin(1), keyserv(1M), newkey(1M), rpc(3NSL), rpc\_clnt\_auth(3NSL), attributes(5)

## send(3SOCKET)

<b>NAME</b>	send, sendto, sendmsg – send a message from a socket				
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lsocket -lnsl [ <i>library</i> ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  ssize_t <b>send</b>(int <i>s</i>, const void *<i>msg</i>, size_t <i>len</i>, int <i>flags</i>); ssize_t <b>sendto</b>(int <i>s</i>, const void *<i>msg</i>, size_t <i>len</i>, int <i>flags</i>, const     struct sockaddr *<i>to</i>, int <i>tolen</i>); ssize_t <b>sendmsg</b>(int <i>s</i>, const struct msghdr *<i>msg</i>, int <i>flags</i>);</pre>				
<b>DESCRIPTION</b>	<p>send(), sendto(), and sendmsg() are used to transmit a message to another transport end-point. send() may be used only when the socket is in a connected state, while sendto() and sendmsg() may be used at any time. <i>s</i> is a socket created with socket(3SOCKET).</p> <p>The address of the target is given by <i>to</i> with <i>tolen</i> specifying its size. The length of the message is given by <i>len</i>. If the message is too long to pass atomically through the underlying protocol, then the error EMSGSIZE is returned, and the message is not transmitted.</p> <p>A return value of -1 indicates locally detected errors only. It does not implicitly mean the message was not delivered.</p> <p>If the socket does not have enough buffer space available to hold the message being sent, send() blocks, unless the socket has been placed in non-blocking I/O mode (seefcntl(2)). The select(3C) or poll(2) call may be used to determine when it is possible to send more data.</p> <p>The <i>flags</i> parameter is formed from the bitwise OR of zero or more of the following:</p> <table><tr><td>MSG_OOB</td><td>Send “out-of-band” data on sockets that support this notion. The underlying protocol must also support “out-of-band” data. Only SOCK_STREAM sockets created in the AF_INET and AF_INET6 address families support out-of-band data.</td></tr><tr><td>MSG_DONTROUTE</td><td>The SO_DONTROUTE option is turned on for the duration of the operation. It is used only by diagnostic or routing programs.</td></tr></table> <p>See recv(3SOCKET) for a description of the msghdr structure.</p>	MSG_OOB	Send “out-of-band” data on sockets that support this notion. The underlying protocol must also support “out-of-band” data. Only SOCK_STREAM sockets created in the AF_INET and AF_INET6 address families support out-of-band data.	MSG_DONTROUTE	The SO_DONTROUTE option is turned on for the duration of the operation. It is used only by diagnostic or routing programs.
MSG_OOB	Send “out-of-band” data on sockets that support this notion. The underlying protocol must also support “out-of-band” data. Only SOCK_STREAM sockets created in the AF_INET and AF_INET6 address families support out-of-band data.				
MSG_DONTROUTE	The SO_DONTROUTE option is turned on for the duration of the operation. It is used only by diagnostic or routing programs.				
<b>RETURN VALUES</b>	These calls return the number of bytes sent, or -1 if an error occurred.				
<b>ERRORS</b>	The send(), sendto(), and sendmsg() functions will fail if:				
EBADF	<i>s</i> is an invalid file descriptor.				
EINTR	The operation was interrupted by delivery of a signal before any data could be buffered to be sent.				



EMSGSIZE	The socket requires that message be sent atomically, and the message was too long.
ENOMEM	There was insufficient memory available to complete the operation.
ENOSR	There were insufficient STREAMS resources available for the operation to complete.
ENOTSOCK	<i>s</i> is not a socket.
EWOULDBLOCK	The socket is marked non-blocking and the requested operation would block.

The `send()` and `sendto()` functions will fail if:

**EINVAL** The *len* argument overflowed a `ssize_t`.

The `sendto()` function will fail if:

**EINVAL** *tolen* is not the size of a valid address for the specified address family.

The `sendmsg()` function will fail if:

**EINVAL** Either the `msg_iovlen` member of the `msg_hdr` structure pointed to by *msg* was less than or equal to 0, or the sum of the *iov\_len* values in the `msg_iov` array overflowed a `ssize_t`.

**EINVAL** Either one of the *iov\_len* values in the `msg_iov` array member of the `msg_hdr` structure pointed to by *msg* was negative, or the sum of the *iov\_len* values in the `msg_iov` array overflowed a `ssize_t`.

The `send()` function will fail if:

**EPIPE** The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, if the socket is of type `SOCK_STREAM`, the `SIGPIPE` signal is generated to the calling process.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `fcntl(2)`, `poll(2)`, `write(2)`, `connect(3SOCKET)`, `getsockopt(3SOCKET)`, `recv(3SOCKET)`, `select(3C)`, `socket(3HEAD)`, `socket(3SOCKET)`, `attributes(5)`

## send(3XNET)

<b>NAME</b>	send – send a message on a socket				
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  ssize_t <b>send</b>(int <i>socket</i>, const void *<i>buffer</i>, size_t <i>length</i>, int <i>flags</i>);</pre>				
<b>DESCRIPTION</b>	<p><i>socket</i> Specifies the socket file descriptor.</p> <p><i>buffer</i> Points to the buffer containing the message to send.</p> <p><i>length</i> Specifies the length of the message in bytes.</p> <p><i>flags</i> Specifies the type of message transmission. Values of this argument are formed by logically OR'ing zero or more of the following flags:</p> <table><tr><td>MSG_EOR</td><td>Terminates a record (if supported by the protocol)</td></tr><tr><td>MSG_OOB</td><td>Sends out-of-band data on sockets that support out-of-band communications. The significance and semantics of out-of-band data are protocol-specific.</td></tr></table>	MSG_EOR	Terminates a record (if supported by the protocol)	MSG_OOB	Sends out-of-band data on sockets that support out-of-band communications. The significance and semantics of out-of-band data are protocol-specific.
MSG_EOR	Terminates a record (if supported by the protocol)				
MSG_OOB	Sends out-of-band data on sockets that support out-of-band communications. The significance and semantics of out-of-band data are protocol-specific.				
	<p>The <code>send()</code> function initiates transmission of a message from the specified socket to its peer. The <code>send()</code> function sends a message only when the socket is connected (including when the peer of a connectionless socket has been set via <code>connect(3XNET)</code>).</p> <p>The length of the message to be sent is specified by the <i>length</i> argument. If the message is too long to pass through the underlying protocol, <code>send()</code> fails and no data is transmitted.</p> <p>Successful completion of a call to <code>send()</code> does not guarantee delivery of the message. A return value of <code>-1</code> indicates only locally-detected errors.</p> <p>If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does not have <code>O_NONBLOCK</code> set, <code>send()</code> blocks until space is available. If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does have <code>O_NONBLOCK</code> set, <code>send()</code> will fail. The <code>select(3C)</code> and <code>poll(2)</code> functions can be used to determine when it is possible to send more data.</p> <p>The socket in use may require the process to have appropriate privileges to use the <code>send()</code> function.</p>				
<b>USAGE</b>	The <code>send()</code> function is identical to <code>sendto(3XNET)</code> with a null pointer <i>dest_len</i> argument, and to <code>write()</code> if no flags are used.				

<b>RETURN VALUES</b>	Upon successful completion, <code>send()</code> returns the number of bytes sent. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.																																		
<b>ERRORS</b>	The <code>send()</code> function will fail if: <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;"><code>EAGAIN</code> <code>EWOULDBLOCK</code></td> <td>The socket's file descriptor is marked <code>O_NONBLOCK</code> and the requested operation would block.</td> </tr> <tr> <td style="vertical-align: top;"><code>EBADF</code></td> <td>The <i>socket</i> argument is not a valid file descriptor.</td> </tr> <tr> <td style="vertical-align: top;"><code>ECONNRESET</code></td> <td>A connection was forcibly closed by a peer.</td> </tr> <tr> <td style="vertical-align: top;"><code>EDESTADDRREQ</code></td> <td>The socket is not connection-mode and no peer address is set.</td> </tr> <tr> <td style="vertical-align: top;"><code>EFAULT</code></td> <td>The <i>buffer</i> parameter can not be accessed.</td> </tr> <tr> <td style="vertical-align: top;"><code>EINTR</code></td> <td>A signal interrupted <code>send()</code> before any data was transmitted.</td> </tr> <tr> <td style="vertical-align: top;"><code>EMSGSIZE</code></td> <td>The message is too large be sent all at once, as the socket requires.</td> </tr> <tr> <td style="vertical-align: top;"><code>ENOTCONN</code></td> <td>The socket is not connected or otherwise has not had the peer prespecified.</td> </tr> <tr> <td style="vertical-align: top;"><code>ENOTSOCK</code></td> <td>The <i>socket</i> argument does not refer to a socket.</td> </tr> <tr> <td style="vertical-align: top;"><code>EOPNOTSUPP</code></td> <td>The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i>.</td> </tr> <tr> <td style="vertical-align: top;"><code>EPIPE</code></td> <td>The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type <code>SOCK_STREAM</code>, the <code>SIGPIPE</code> signal is generated to the calling process.</td> </tr> </table> <p>The <code>send()</code> function may fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;"><code>EACCES</code></td> <td>The calling process does not have the appropriate privileges.</td> </tr> <tr> <td style="vertical-align: top;"><code>EIO</code></td> <td>An I/O error occurred while reading from or writing to the file system.</td> </tr> <tr> <td style="vertical-align: top;"><code>ENETDOWN</code></td> <td>The local interface used to reach the destination is down.</td> </tr> <tr> <td style="vertical-align: top;"><code>ENETUNREACH</code></td> <td>No route to the network is present.</td> </tr> <tr> <td style="vertical-align: top;"><code>ENOBUFS</code></td> <td>Insufficient resources were available in the system to perform the operation.</td> </tr> <tr> <td style="vertical-align: top;"><code>ENOSR</code></td> <td>There were insufficient <code>STREAMS</code> resources available for the operation to complete.</td> </tr> </table>	<code>EAGAIN</code> <code>EWOULDBLOCK</code>	The socket's file descriptor is marked <code>O_NONBLOCK</code> and the requested operation would block.	<code>EBADF</code>	The <i>socket</i> argument is not a valid file descriptor.	<code>ECONNRESET</code>	A connection was forcibly closed by a peer.	<code>EDESTADDRREQ</code>	The socket is not connection-mode and no peer address is set.	<code>EFAULT</code>	The <i>buffer</i> parameter can not be accessed.	<code>EINTR</code>	A signal interrupted <code>send()</code> before any data was transmitted.	<code>EMSGSIZE</code>	The message is too large be sent all at once, as the socket requires.	<code>ENOTCONN</code>	The socket is not connected or otherwise has not had the peer prespecified.	<code>ENOTSOCK</code>	The <i>socket</i> argument does not refer to a socket.	<code>EOPNOTSUPP</code>	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .	<code>EPIPE</code>	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type <code>SOCK_STREAM</code> , the <code>SIGPIPE</code> signal is generated to the calling process.	<code>EACCES</code>	The calling process does not have the appropriate privileges.	<code>EIO</code>	An I/O error occurred while reading from or writing to the file system.	<code>ENETDOWN</code>	The local interface used to reach the destination is down.	<code>ENETUNREACH</code>	No route to the network is present.	<code>ENOBUFS</code>	Insufficient resources were available in the system to perform the operation.	<code>ENOSR</code>	There were insufficient <code>STREAMS</code> resources available for the operation to complete.
<code>EAGAIN</code> <code>EWOULDBLOCK</code>	The socket's file descriptor is marked <code>O_NONBLOCK</code> and the requested operation would block.																																		
<code>EBADF</code>	The <i>socket</i> argument is not a valid file descriptor.																																		
<code>ECONNRESET</code>	A connection was forcibly closed by a peer.																																		
<code>EDESTADDRREQ</code>	The socket is not connection-mode and no peer address is set.																																		
<code>EFAULT</code>	The <i>buffer</i> parameter can not be accessed.																																		
<code>EINTR</code>	A signal interrupted <code>send()</code> before any data was transmitted.																																		
<code>EMSGSIZE</code>	The message is too large be sent all at once, as the socket requires.																																		
<code>ENOTCONN</code>	The socket is not connected or otherwise has not had the peer prespecified.																																		
<code>ENOTSOCK</code>	The <i>socket</i> argument does not refer to a socket.																																		
<code>EOPNOTSUPP</code>	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .																																		
<code>EPIPE</code>	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type <code>SOCK_STREAM</code> , the <code>SIGPIPE</code> signal is generated to the calling process.																																		
<code>EACCES</code>	The calling process does not have the appropriate privileges.																																		
<code>EIO</code>	An I/O error occurred while reading from or writing to the file system.																																		
<code>ENETDOWN</code>	The local interface used to reach the destination is down.																																		
<code>ENETUNREACH</code>	No route to the network is present.																																		
<code>ENOBUFS</code>	Insufficient resources were available in the system to perform the operation.																																		
<code>ENOSR</code>	There were insufficient <code>STREAMS</code> resources available for the operation to complete.																																		

send(3XNET)

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `connect(3XNET)`, `getsockopt(3XNET)`, `poll(2)`, `recv(3XNET)`, `recvfrom(3XNET)`, `recvmsg(3XNET)`, `select(3C)`, `sendmsg(3XNET)`, `sendto(3XNET)`, `setsockopt(3XNET)`, `shutdown(3XNET)`, `socket(3XNET)`, `attributes(5)`

<b>NAME</b>	sendmsg – send a message on a socket using a message structure										
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  ssize_t <b>sendmsg</b>(int <i>socket</i>, const struct msghdr *<i>message</i>, int <i>flags</i>);</pre>										
<b>DESCRIPTION</b>	<p>The <code>sendmsg()</code> function sends a message through a connection-mode or connectionless-mode socket. If the socket is connectionless-mode, the message will be sent to the address specified by <code>msghdr</code>. If the socket is connection-mode, the destination address in <code>msghdr</code> is ignored.</p> <p>The function takes the following arguments:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><i>socket</i></td> <td>Specifies the socket file descriptor.</td> </tr> <tr> <td><i>message</i></td> <td>Points to a <code>msghdr</code> structure, containing both the destination address and the buffers for the outgoing message. The length and format of the address depend on the address family of the socket. The <code>msg_flags</code> member is ignored.</td> </tr> <tr> <td><i>flags</i></td> <td>Specifies the type of message transmission. The application may specify 0 or the following flag:</td> </tr> <tr> <td style="padding-left: 40px;">MSG_EOR</td> <td>Terminates a record (if supported by the protocol)</td> </tr> <tr> <td style="padding-left: 40px;">MSG_OOB</td> <td>Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.</td> </tr> </table> <p>The <code>msg_iov</code> and <code>msg_iovlen</code> fields of message specify zero or more buffers containing the data to be sent. <code>msg_iov</code> points to an array of <code>iovec</code> structures; <code>msg_iovlen</code> must be set to the dimension of this array. In each <code>iovec</code> structure, the <code>iov_base</code> field specifies a storage area and the <code>iov_len</code> field gives its size in bytes. Some of these sizes can be zero. The data from each storage area indicated by <code>msg_iov</code> is sent in turn.</p> <p>Successful completion of a call to <code>sendmsg()</code> does not guarantee delivery of the message. A return value of <code>-1</code> indicates only locally-detected errors.</p> <p>If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does not have <code>O_NONBLOCK</code> set, <code>sendmsg()</code> function blocks until space is available. If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does have <code>O_NONBLOCK</code> set, <code>sendmsg()</code> function will fail.</p> <p>If the socket protocol supports broadcast and the specified address is a broadcast address for the socket protocol, <code>sendmsg()</code> will fail if the <code>SO_BROADCAST</code> option is not set for the socket.</p>	<i>socket</i>	Specifies the socket file descriptor.	<i>message</i>	Points to a <code>msghdr</code> structure, containing both the destination address and the buffers for the outgoing message. The length and format of the address depend on the address family of the socket. The <code>msg_flags</code> member is ignored.	<i>flags</i>	Specifies the type of message transmission. The application may specify 0 or the following flag:	MSG_EOR	Terminates a record (if supported by the protocol)	MSG_OOB	Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.
<i>socket</i>	Specifies the socket file descriptor.										
<i>message</i>	Points to a <code>msghdr</code> structure, containing both the destination address and the buffers for the outgoing message. The length and format of the address depend on the address family of the socket. The <code>msg_flags</code> member is ignored.										
<i>flags</i>	Specifies the type of message transmission. The application may specify 0 or the following flag:										
MSG_EOR	Terminates a record (if supported by the protocol)										
MSG_OOB	Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.										

## sendmsg(3XNET)

	The socket in use may require the process to have appropriate privileges to use the <code>sendmsg()</code> function.																										
<b>USAGE</b>	The <code>select(3C)</code> and <code>poll(2)</code> functions can be used to determine when it is possible to send more data.																										
<b>RETURN VALUES</b>	Upon successful completion, <code>sendmsg()</code> function returns the number of bytes sent. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.																										
<b>ERRORS</b>	The <code>sendmsg()</code> function will fail if:  <table><tr><td><code>EAGAIN</code></td><td></td></tr><tr><td><code>EWOULDBLOCK</code></td><td>The socket's file descriptor is marked <code>O_NONBLOCK</code> and the requested operation would block.</td></tr><tr><td><code>EAFNOSUPPORT</code></td><td>Addresses in the specified address family cannot be used with this socket.</td></tr><tr><td><code>EBADF</code></td><td>The <i>socket</i> argument is not a valid file descriptor.</td></tr><tr><td><code>ECONNRESET</code></td><td>A connection was forcibly closed by a peer.</td></tr><tr><td><code>EFAULT</code></td><td>The <i>message</i> parameter, or storage pointed to by the <i>msg_name</i>, <i>msg_control</i> or <i>msg_iov</i> fields of the <i>message</i> parameter, or storage pointed to by the <i>iovec</i> structures pointed to by the <i>msg_iov</i> field can not be accessed.</td></tr><tr><td><code>EINTR</code></td><td>A signal interrupted <code>sendmsg()</code> before any data was transmitted.</td></tr><tr><td><code>EINVAL</code></td><td>The sum of the <i>iov_len</i> values overflows an <code>ssize_t</code>.</td></tr><tr><td><code>EMSGSIZE</code></td><td>The message is too large to be sent all at once (as the socket requires), or the <i>msg_iovlen</i> member of the <i>msg_hdr</i> structure pointed to by <i>message</i> is less than or equal to 0 or is greater than <code>IOV_MAX</code>.</td></tr><tr><td><code>ENOTCONN</code></td><td>The socket is connection-mode but is not connected.</td></tr><tr><td><code>ENOTSOCK</code></td><td>The <i>socket</i> argument does not refer a socket.</td></tr><tr><td><code>EOPNOTSUPP</code></td><td>The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i>.</td></tr><tr><td><code>EPIPE</code></td><td>The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type <code>SOCK_STREAM</code>, the <code>SIGPIPE</code> signal is generated to the calling process.</td></tr></table> If the address family of the socket is <code>AF_UNIX</code> , then <code>sendmsg()</code> will fail if:	<code>EAGAIN</code>		<code>EWOULDBLOCK</code>	The socket's file descriptor is marked <code>O_NONBLOCK</code> and the requested operation would block.	<code>EAFNOSUPPORT</code>	Addresses in the specified address family cannot be used with this socket.	<code>EBADF</code>	The <i>socket</i> argument is not a valid file descriptor.	<code>ECONNRESET</code>	A connection was forcibly closed by a peer.	<code>EFAULT</code>	The <i>message</i> parameter, or storage pointed to by the <i>msg_name</i> , <i>msg_control</i> or <i>msg_iov</i> fields of the <i>message</i> parameter, or storage pointed to by the <i>iovec</i> structures pointed to by the <i>msg_iov</i> field can not be accessed.	<code>EINTR</code>	A signal interrupted <code>sendmsg()</code> before any data was transmitted.	<code>EINVAL</code>	The sum of the <i>iov_len</i> values overflows an <code>ssize_t</code> .	<code>EMSGSIZE</code>	The message is too large to be sent all at once (as the socket requires), or the <i>msg_iovlen</i> member of the <i>msg_hdr</i> structure pointed to by <i>message</i> is less than or equal to 0 or is greater than <code>IOV_MAX</code> .	<code>ENOTCONN</code>	The socket is connection-mode but is not connected.	<code>ENOTSOCK</code>	The <i>socket</i> argument does not refer a socket.	<code>EOPNOTSUPP</code>	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .	<code>EPIPE</code>	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type <code>SOCK_STREAM</code> , the <code>SIGPIPE</code> signal is generated to the calling process.
<code>EAGAIN</code>																											
<code>EWOULDBLOCK</code>	The socket's file descriptor is marked <code>O_NONBLOCK</code> and the requested operation would block.																										
<code>EAFNOSUPPORT</code>	Addresses in the specified address family cannot be used with this socket.																										
<code>EBADF</code>	The <i>socket</i> argument is not a valid file descriptor.																										
<code>ECONNRESET</code>	A connection was forcibly closed by a peer.																										
<code>EFAULT</code>	The <i>message</i> parameter, or storage pointed to by the <i>msg_name</i> , <i>msg_control</i> or <i>msg_iov</i> fields of the <i>message</i> parameter, or storage pointed to by the <i>iovec</i> structures pointed to by the <i>msg_iov</i> field can not be accessed.																										
<code>EINTR</code>	A signal interrupted <code>sendmsg()</code> before any data was transmitted.																										
<code>EINVAL</code>	The sum of the <i>iov_len</i> values overflows an <code>ssize_t</code> .																										
<code>EMSGSIZE</code>	The message is too large to be sent all at once (as the socket requires), or the <i>msg_iovlen</i> member of the <i>msg_hdr</i> structure pointed to by <i>message</i> is less than or equal to 0 or is greater than <code>IOV_MAX</code> .																										
<code>ENOTCONN</code>	The socket is connection-mode but is not connected.																										
<code>ENOTSOCK</code>	The <i>socket</i> argument does not refer a socket.																										
<code>EOPNOTSUPP</code>	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .																										
<code>EPIPE</code>	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type <code>SOCK_STREAM</code> , the <code>SIGPIPE</code> signal is generated to the calling process.																										

EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating the pathname in the socket address.
ENAMETOOLONG	A component of a pathname exceeded NAME_MAX characters, or an entire pathname exceeded PATH_MAX characters.
ENOENT	A component of the pathname does not name an existing file or the pathname is an empty string.
ENOTDIR	A component of the path prefix of the pathname in the socket address is not a directory.

The sendmsg ( ) function may fail if:

EACCES	Search permission is denied for a component of the path prefix; or write access to the named socket is denied.
EDESTADDRREQ	The socket is not connection-mode and does not have its peer address set, and no destination address was specified.
EHOSTUNREACH	The destination host cannot be reached (probably because the host is down or a remote router cannot reach it).
EIO	An I/O error occurred while reading from or writing to the file system.
EISCONN	A destination address was specified and the socket is already connected.
ENETDOWN	The local interface used to reach the destination is down.
ENETUNREACH	No route to the network is present.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
ENOMEM	Insufficient memory was available to fulfill the request.
ENOSR	There were insufficient STREAMS resources available for the operation to complete.

If the address family of the socket is AF\_UNIX, then sendmsg ( ) may fail if:

ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds PATH_MAX.
--------------	---

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

sendmsg(3XNET)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** poll(2) getsockopt(3XNET), recv(3XNET), recvfrom(3XNET),  
recvmsg(3XNET), select(3C), send(3XNET), sendto(3XNET),  
setsockopt(3XNET), shutdown(3XNET), socket(3XNET), attributes(5)



<b>NAME</b>	sendto – send a message on a socket												
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lxnet [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;</pre> <p> <code>ssize_t <b>sendto</b>(int <i>socket</i>, const void *<i>message</i>, size_t <i>length</i>, int <i>flags</i>, const struct sockaddr *<i>dest_addr</i>, socklen_t <i>dest_len</i>);</code> </p>												
<b>DESCRIPTION</b>	<p>The <code>sendto()</code> function sends a message through a connection-mode or connectionless-mode socket. If the socket is connectionless-mode, the message will be sent to the address specified by <i>dest_addr</i>. If the socket is connection-mode, <i>dest_addr</i> is ignored.</p> <p>The function takes the following arguments:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><i>socket</i></td> <td>Specifies the socket file descriptor.</td> </tr> <tr> <td><i>message</i></td> <td>Points to a buffer containing the message to be sent.</td> </tr> <tr> <td><i>length</i></td> <td>Specifies the size of the message in bytes.</td> </tr> <tr> <td><i>flags</i></td> <td>Specifies the type of message transmission. Values of this argument are formed by logically OR'ing zero or more of the following flags:</td> </tr> <tr> <td style="padding-left: 40px;">MSG_EOR</td> <td>Terminates a record (if supported by the protocol)</td> </tr> <tr> <td style="padding-left: 40px;">MSG_OOB</td> <td>Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.</td> </tr> </table> <p><i>dest_addr</i> Points to a <code>sockaddr</code> structure containing the destination address. The length and format of the address depend on the address family of the socket.</p> <p><i>dest_len</i> Specifies the length of the <code>sockaddr</code> structure pointed to by the <i>dest_addr</i> argument.</p> <p>If the socket protocol supports broadcast and the specified address is a broadcast address for the socket protocol, <code>sendto()</code> will fail if the <code>SO_BROADCAST</code> option is not set for the socket.</p> <p>The <i>dest_addr</i> argument specifies the address of the target. The <i>length</i> argument specifies the length of the message.</p> <p>Successful completion of a call to <code>sendto()</code> does not guarantee delivery of the message. A return value of <code>-1</code> indicates only locally-detected errors.</p>	<i>socket</i>	Specifies the socket file descriptor.	<i>message</i>	Points to a buffer containing the message to be sent.	<i>length</i>	Specifies the size of the message in bytes.	<i>flags</i>	Specifies the type of message transmission. Values of this argument are formed by logically OR'ing zero or more of the following flags:	MSG_EOR	Terminates a record (if supported by the protocol)	MSG_OOB	Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.
<i>socket</i>	Specifies the socket file descriptor.												
<i>message</i>	Points to a buffer containing the message to be sent.												
<i>length</i>	Specifies the size of the message in bytes.												
<i>flags</i>	Specifies the type of message transmission. Values of this argument are formed by logically OR'ing zero or more of the following flags:												
MSG_EOR	Terminates a record (if supported by the protocol)												
MSG_OOB	Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.												

## sendto(3XNET)

If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does not have `O_NONBLOCK` set, `sendto()` blocks until space is available. If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does have `O_NONBLOCK` set, `sendto()` will fail.

The socket in use may require the process to have appropriate privileges to use the `sendto()` function.

**USAGE** The `select(3C)` and `poll(2)` functions can be used to determine when it is possible to send more data.

**RETURN VALUES** Upon successful completion, `sendto()` returns the number of bytes sent. Otherwise, `-1` is returned and `errno` is set to indicate the error.

**ERRORS** The `sendto()` function will fail if:

<code>EAFNOSUPPORT</code>	Addresses in the specified address family cannot be used with this socket.
<code>EAGAIN</code> <code>EWOULDBLOCK</code>	The socket's file descriptor is marked <code>O_NONBLOCK</code> and the requested operation would block.
<code>EBADF</code>	The <i>socket</i> argument is not a valid file descriptor.
<code>ECONNRESET</code>	A connection was forcibly closed by a peer.
<code>EFAULT</code>	The <i>message</i> or <i>destaddr</i> parameter can not be accessed.
<code>EINTR</code>	A signal interrupted <code>sendto()</code> before any data was transmitted.
<code>EMSGSIZE</code>	The message is too large to be sent all at once, as the socket requires.
<code>ENOTCONN</code>	The socket is connection-mode but is not connected.
<code>ENOTSOCK</code>	The <i>socket</i> argument does not refer to a socket.
<code>EOPNOTSUPP</code>	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .
<code>EPIPE</code>	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type <code>SOCK_STREAM</code> , the <code>SIGPIPE</code> signal is generated to the calling process.

If the address family of the socket is `AF_UNIX`, then `sendto()` will fail if:

<code>EIO</code>	An I/O error occurred while reading from or writing to the file system.
------------------	---

ELOOP	Too many symbolic links were encountered in translating the pathname in the socket address.
ENAMETOOLONG	A component of a pathname exceeded <code>NAME_MAX</code> characters, or an entire pathname exceeded <code>PATH_MAX</code> characters.
ENOENT	A component of the pathname does not name an existing file or the pathname is an empty string.
ENOTDIR	A component of the path prefix of the pathname in the socket address is not a directory.

The `sendto()` function may fail if:

EACCES	Search permission is denied for a component of the path prefix; or write access to the named socket is denied.
EDESTADDRREQ	The socket is not connection-mode and does not have its peer address set, and no destination address was specified.
EHOSTUNREACH	The destination host cannot be reached (probably because the host is down or a remote router cannot reach it).
EINVAL	The <code>dest_len</code> argument is not a valid length for the address family.
EIO	An I/O error occurred while reading from or writing to the file system.
EISCONN	A destination address was specified and the socket is already connected.
ENETDOWN	The local interface used to reach the destination is down.
ENETUNREACH	No route to the network is present.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
ENOMEM	Insufficient memory was available to fulfill the request.
ENOSR	There were insufficient STREAMS resources available for the operation to complete.

If the address family of the socket is `AF_UNIX`, then `sendto()` may fail if:

ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> .
--------------	---

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

sendto(3XNET)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** poll(2), getsockopt(3XNET), recv(3XNET), recvfrom(3XNET),  
recvmsg(3XNET), select(3C), send(3XNET), sendmsg(3XNET),  
setsockopt(3XNET), shutdown(3XNET), socket(3XNET), attributes(5)

<b>NAME</b>	setsockopt – set the socket options														
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;sys/socket.h&gt;  int <b>setsockopt</b>(int <i>socket</i>, int <i>level</i>, int <i>option_name</i>, const                 void*<i>option_value</i>, socklen_t <i>option_len</i>);</pre>														
<b>DESCRIPTION</b>	<p>The <code>setsockopt()</code> function sets the option specified by the <code>option_name</code> argument, at the protocol level specified by the <code>level</code> argument, to the value pointed to by the <code>option_value</code> argument for the socket associated with the file descriptor specified by the <code>socket</code> argument.</p> <p>The <code>level</code> argument specifies the protocol level at which the option resides. To set options at the socket level, specify the <code>level</code> argument as <code>SOL_SOCKET</code>. To set options at other levels, supply the appropriate protocol number for the protocol controlling the option. For example, to indicate that an option will be interpreted by the TCP (Transport Control Protocol), set <code>level</code> to the protocol number of TCP, as defined in the <code>&lt;netinet/in.h&gt;</code> header, or as determined by using <code>getprotobyname(3XNET)</code>.</p> <p>The <code>option_name</code> argument specifies a single option to set. The <code>option_name</code> argument and any specified options are passed uninterpreted to the appropriate protocol module for interpretations. The <code>&lt;sys/socket.h&gt;</code> header defines the socket level options. The options are as follows:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">SO_DEBUG</td> <td>Turns on recording of debugging information. This option enables or disables debugging in the underlying protocol modules. This option takes an <code>int</code> value. This is a boolean option.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">SO_BROADCAST</td> <td>Permits sending of broadcast messages, if this is supported by the protocol. This option takes an <code>int</code> value. This is a boolean option.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">SO_REUSEADDR</td> <td>Specifies that the rules used in validating addresses supplied to <code>bind(3XNET)</code> should allow reuse of local addresses, if this is supported by the protocol. This option takes an <code>int</code> value. This is a boolean option.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">SO_KEEPALIVE</td> <td>Keeps connections active by enabling the periodic transmission of messages, if this is supported by the protocol. This option takes an <code>int</code> value.</td> </tr> <tr> <td></td> <td>If the connected socket fails to respond to these messages, the connection is broken and processes writing to that socket are notified with a <code>SIGPIPE</code> signal.</td> </tr> <tr> <td></td> <td>This is a boolean option.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">SO_LINGER</td> <td>Lingers on a <code>close(2)</code> if data is present. This option controls the action taken when unsent messages queue</td> </tr> </table>	SO_DEBUG	Turns on recording of debugging information. This option enables or disables debugging in the underlying protocol modules. This option takes an <code>int</code> value. This is a boolean option.	SO_BROADCAST	Permits sending of broadcast messages, if this is supported by the protocol. This option takes an <code>int</code> value. This is a boolean option.	SO_REUSEADDR	Specifies that the rules used in validating addresses supplied to <code>bind(3XNET)</code> should allow reuse of local addresses, if this is supported by the protocol. This option takes an <code>int</code> value. This is a boolean option.	SO_KEEPALIVE	Keeps connections active by enabling the periodic transmission of messages, if this is supported by the protocol. This option takes an <code>int</code> value.		If the connected socket fails to respond to these messages, the connection is broken and processes writing to that socket are notified with a <code>SIGPIPE</code> signal.		This is a boolean option.	SO_LINGER	Lingers on a <code>close(2)</code> if data is present. This option controls the action taken when unsent messages queue
SO_DEBUG	Turns on recording of debugging information. This option enables or disables debugging in the underlying protocol modules. This option takes an <code>int</code> value. This is a boolean option.														
SO_BROADCAST	Permits sending of broadcast messages, if this is supported by the protocol. This option takes an <code>int</code> value. This is a boolean option.														
SO_REUSEADDR	Specifies that the rules used in validating addresses supplied to <code>bind(3XNET)</code> should allow reuse of local addresses, if this is supported by the protocol. This option takes an <code>int</code> value. This is a boolean option.														
SO_KEEPALIVE	Keeps connections active by enabling the periodic transmission of messages, if this is supported by the protocol. This option takes an <code>int</code> value.														
	If the connected socket fails to respond to these messages, the connection is broken and processes writing to that socket are notified with a <code>SIGPIPE</code> signal.														
	This is a boolean option.														
SO_LINGER	Lingers on a <code>close(2)</code> if data is present. This option controls the action taken when unsent messages queue														

## setsockopt(3XNET)

	<p>on a socket and <code>close(2)</code> is performed. If <code>SO_LINGER</code> is set, the system blocks the process during <code>close(2)</code> until it can transmit the data or until the time expires. If <code>SO_LINGER</code> is not specified, and <code>close(2)</code> is issued, the system handles the call in a way that allows the process to continue as quickly as possible. This option takes a <code>linger</code> structure, as defined in the <code>&lt;sys/socket.h&gt;</code> header, to specify the state of the option and linger interval.</p>
	<p><code>SO_OOBINLINE</code> Leaves received out-of-band data (data marked urgent) in line. This option takes an <code>int</code> value. This is a boolean option.</p>
	<p><code>SO_SNDBUF</code> Sets send buffer size. This option takes an <code>int</code> value.</p>
	<p><code>SO_RCVBUF</code> Sets receive buffer size. This option takes an <code>int</code> value.</p>
	<p><code>SO_DONTROUTE</code> Requests that outgoing messages bypass the standard routing facilities. The destination must be on a directly-connected network, and messages are directed to the appropriate network interface according to the destination address. The effect, if any, of this option depends on what protocol is in use. This option takes an <code>int</code> value. This is a boolean option.</p>
	<p>For boolean options, 0 indicates that the option is disabled and 1 indicates that the option is enabled.</p>
	<p>Options at other protocol levels vary in format and name.</p>
<b>USAGE</b>	<p>The <code>setsockopt()</code> function provides an application program with the means to control socket behavior. An application program can use <code>setsockopt()</code> to allocate buffer space, control timeouts, or permit socket data broadcasts. The <code>&lt;sys/socket.h&gt;</code> header defines the socket-level options available to <code>setsockopt()</code>.</p>
	<p>Options may exist at multiple protocol levels. The <code>SO_</code> options are always present at the uppermost socket level.</p>
<b>RETURN VALUES</b>	<p>Upon successful completion, <code>setsockopt()</code> returns 0. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.</p>
<b>ERRORS</b>	<p>The <code>setsockopt()</code> function will fail if:</p>
	<p><code>EBADF</code> The <i>socket</i> argument is not a valid file descriptor.</p>
	<p><code>EDOM</code> The send and receive timeout values are too big to fit into the timeout fields in the socket structure.</p>
	<p><code>EFAULT</code> The <i>option_value</i> parameter can not be accessed or written.</p>

setsockopt(3XNET)

- EINVAL            The specified option is invalid at the specified socket level or the socket has been shut down.
- EISCONN         The socket is already connected, and a specified option can not be set while the socket is connected.
- ENOPROTOOPT    The option is not supported by the protocol.
- ENOTSOCK        The *socket* argument does not refer to a socket.

The `setsockopt()` function may fail if:

- ENOMEM           There was insufficient memory available for the operation to complete.
- ENOBUFS          Insufficient resources are available in the system to complete the call.
- ENOSR            There were insufficient STREAMS resources available for the operation to complete.

**ATTRIBUTES**    See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO**    `bind(3XNET)`, `endprotoent(3XNET)`, `getsockopt(3XNET)`, `socket(3XNET)`, `attributes(5)`

## shutdown(3SOCKET)

**NAME** shutdown – shut down part of a full-duplex connection

**SYNOPSIS** `cc [ flag ... ] file ... -lsocket -lnsl [ library ... ]`  
`int shutdown(int s, int how);`

**DESCRIPTION** The `shutdown()` call shuts down all or part of a full-duplex connection on the socket associated with `s`. If `how` is 0, then further receives will be disallowed. If `how` is 1, then further sends will be disallowed. If `how` is 2, then further sends and receives will be disallowed.

**RETURN VALUES** A 0 is returned if the call succeeds, -1 if it fails.

**ERRORS** The call succeeds unless:

EBADF	<code>s</code> is not a valid file descriptor.
ENOMEM	There was insufficient user memory available for the operation to complete.
ENOSR	There were insufficient STREAMS resources available for the operation to complete.
ENOTCONN	The specified socket is not connected.
ENOTSOCK	<code>s</code> is not a socket.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `connect(3SOCKET)`, `socket(3SOCKET)`, `attributes(5)`, `socket(3HEAD)`

**NOTES** The `how` values should be defined constants.



<b>NAME</b>	shutdown – shut down socket send and receive operations												
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;sys/socket.h&gt;  int <b>shutdown</b>(int <i>socket</i>, int <i>how</i>);</pre>												
<b>DESCRIPTION</b>	<p><i>socket</i> Specifies the file descriptor of the socket.</p> <p><i>how</i> Specifies the type of shutdown. The values are as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;">SHUT_RD</td> <td>Disables further receive operations.</td> </tr> <tr> <td style="padding-right: 20px;">SHUT_WR</td> <td>Disables further send operations.</td> </tr> <tr> <td style="padding-right: 20px;">SHUT_RDWR</td> <td>Disables further send and receive operations.</td> </tr> </table> <p>The <code>shutdown()</code> function disables subsequent send and/or receive operations on a socket, depending on the value of the <i>how</i> argument.</p>	SHUT_RD	Disables further receive operations.	SHUT_WR	Disables further send operations.	SHUT_RDWR	Disables further send and receive operations.						
SHUT_RD	Disables further receive operations.												
SHUT_WR	Disables further send operations.												
SHUT_RDWR	Disables further send and receive operations.												
<b>RETURN VALUES</b>	Upon successful completion, <code>shutdown()</code> returns 0. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.												
<b>ERRORS</b>	<p>The <code>shutdown()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EBADF</td> <td>The <i>socket</i> argument is not a valid file descriptor.</td> </tr> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The <i>how</i> argument is invalid.</td> </tr> <tr> <td style="padding-right: 20px;">ENOTCONN</td> <td>The socket is not connected.</td> </tr> <tr> <td style="padding-right: 20px;">ENOTSOCK</td> <td>The <i>socket</i> argument does not refer to a socket.</td> </tr> </table> <p>The <code>shutdown()</code> function may fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">ENOBUFS</td> <td>Insufficient resources were available in the system to perform the operation.</td> </tr> <tr> <td style="padding-right: 20px;">ENOSR</td> <td>There were insufficient STREAMS resources available for the operation to complete.</td> </tr> </table>	EBADF	The <i>socket</i> argument is not a valid file descriptor.	EINVAL	The <i>how</i> argument is invalid.	ENOTCONN	The socket is not connected.	ENOTSOCK	The <i>socket</i> argument does not refer to a socket.	ENOBUFS	Insufficient resources were available in the system to perform the operation.	ENOSR	There were insufficient STREAMS resources available for the operation to complete.
EBADF	The <i>socket</i> argument is not a valid file descriptor.												
EINVAL	The <i>how</i> argument is invalid.												
ENOTCONN	The socket is not connected.												
ENOTSOCK	The <i>socket</i> argument does not refer to a socket.												
ENOBUFS	Insufficient resources were available in the system to perform the operation.												
ENOSR	There were insufficient STREAMS resources available for the operation to complete.												
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:												
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe								
ATTRIBUTE TYPE	ATTRIBUTE VALUE												
MT-Level	MT-Safe												
<b>SEE ALSO</b>	<code>getsockopt(3XNET)</code> , <code>recv(3XNET)</code> , <code>recvfrom(3XNET)</code> , <code>recvmsg(3XNET)</code> , <code>select(3C)</code> , <code>send(3XNET)</code> , <code>sendto(3XNET)</code> , <code>setsockopt(3XNET)</code> , <code>socket(3XNET)</code> , <code>attributes(5)</code>												

## slp\_api(3SLP)

<b>NAME</b>	slp_api – Service Location Protocol Application Programming Interface
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lslp [ library ... ] #include &lt;slp.h&gt;</pre>
<b>DESCRIPTION</b>	<p>The <code>slp_api</code> is a C language binding that maps directly into the Service Location Protocol (“SLP”) defined by <i>RFC 2614</i>. This implementation requires minimal overhead. With the exception of the <code>SLPDereg()</code> and <code>SLPDelAttrs()</code> functions, which map into different uses of the SLP deregister request, there is one C language function per protocol request. Parameters are for the most part character buffers. Memory management is kept simple because the client allocates most memory and client callback functions are required to copy incoming parameters into memory allocated by the client code. Any memory returned directly from the API functions is deallocated using the <code>SLPFree()</code> function.</p> <p>To conform with standard C practice, all character strings passed to and returned through the API are null-terminated, even though the SLP protocol does not use null-terminated strings. Strings passed as parameters are UTF-8 but they may still be passed as a C string (a null-terminated sequence of bytes.) Escaped characters must be encoded by the API client as UTF-8. In the common case of US-ASCII, the usual one byte per character C strings work. API functions assist in escaping and unescaping strings.</p> <p>Unless otherwise noted, parameters to API functions and callbacks are non-NULL. Some parameters may have other restrictions. If any parameter fails to satisfy the restrictions on its value, the operation returns a <code>PARAMETER_BAD</code> error.</p>
<b>Syntax for String Parameters</b>	<p>Query strings, attribute registration lists, attribute deregistration lists, scope lists, and attribute selection lists follow the syntax described in <i>RFC 2608</i>. The API reflects the strings passed from clients directly into protocol requests, and reflects out strings returned from protocol replies directly to clients. As a consequence, clients are responsible for formatting request strings, including escaping and converting opaque values to escaped byte-encoded strings. Similarly, on output, clients are required to unescape strings and convert escaped string-encoded opaques to binary. The <code>SLPEscape()</code> and <code>SLPUnescape()</code> functions can be used for escaping SLP reserved characters, but they perform no opaque processing.</p> <p>Opaque values consist of a character buffer that contains a UTF-8-encoded string, the first characters of which are the non UTF-8 encoding “\xff”. Subsequent characters are the escaped values for the original bytes in the opaque. The escape convention is relatively simple. An escape consists of a backslash followed by the two hexadecimal digits encoding the byte. An example is “\2c” for the byte <code>0x2c</code>. Clients handle opaque processing themselves, since the algorithm is relatively simple and uniform.</p>
<b>System Properties</b>	<p>The system properties established in <code>slp.conf(4)</code>, the configuration file, are accessible through the <code>SLPGetProperty()</code> and <code>SLPSetProperty()</code> functions. The <code>SLPSetProperty()</code> function modifies properties only in the running process, not in the configuration file. Errors are checked when the property is used and, as with parsing the configuration file, are logged at the <code>LOG_INFO</code> priority. Program execution continues without interruption by substituting the default for the erroneous</p>

parameter. In general, individual agents should rarely be required to override these properties, since they reflect properties of the SLP network that are not of concern to individual agents. If changes are required, system administrators should modify the configuration file.

Properties are global to the process, affecting all threads and all handles created with `SLPOpen()`.

### Memory Management

The only API functions that return memory specifically requiring deallocation on the part of the client are `SLPParseSrvURL()`, `SLPFindScope()`, `SLPEscape()`, and `SLPUnescape()`. Free this memory with `SLPFree()` when it is no longer needed. Do not free character strings returned by means of the `SLPGetProperty()` function.

Any memory passed to callbacks belongs to the library, and it must not be retained by the client code. Otherwise, crashes are possible. Clients must copy data out of the callback parameters. No other use of the memory in callback parameters is allowed.

### Asynchronous and Incremental Return Semantics

If a handle parameter to an API function is opened asynchronously, the API function calls on the handle to check the other parameters, opens the appropriate operation, and returns immediately. If an error occurs in the process of starting the operation, the error code is returned. If the handle parameter is opened synchronously, the function call is blocked until all results are available, and it returns only after the results are reported through the callback function. The return code indicates whether any errors occurred during the operation.

The callback function is called whenever the API library has results to report. The callback code is required to check the error code parameter before looking at the other parameters. If the error code is not `SLP_OK`, the other parameters may be `NULL` or otherwise invalid. The API library can terminate any outstanding operation on which an error occurs. The callback code can similarly indicate that the operation should be terminated by passing back `SLP_FALSE` to indicate that it is not interested in receiving more results. Callback functions are not permitted to recursively call into the API on the same `SLPHandle`. If an attempt is made to call into the API, the API function returns `SLP_HANDLE_IN_USE`. Prohibiting recursive callbacks on the same handle simplifies implementation of thread safe code, since locks held on the handle will not be in place during a second outcall on the handle.

The total number of results received can be controlled by setting the `net.slp.maxResults` parameter.

On the last call to a callback, whether asynchronous or synchronous, the status code passed to the callback has value `SLP_LAST_CALL`. There are four reasons why the call can terminate:

DA reply received

A reply from a DA has been received and therefore nothing more is expected.

Multicast terminated

The multicast convergence time has elapsed and the API library multicast code is giving up.

## slp\_api(3SLP)

	Multicast null results	Nothing new has been received during multicast for awhile and the API library multicast code is giving up on that (as an optimization).
	Maximum results	The user has set the <code>net.slp.maxResults</code> property and that number of replies has been collected and returned.
<b>Configuration Files</b>	The API library reads <code>slp.conf(4)</code> , the default configuration file, to obtain the operating parameters. You can specify the location of this file with the <code>SLP_CONF_FILE</code> environment variable. If you do not set this variable, or the file it refers to is invalid, the API will use the default configuration file at <code>/etc/inet/slp.conf</code> instead.	
<b>Data Structures</b>	The data structures used by the SLP API are as follows:  <b>The URL Lifetime Type</b>  <pre>typedef enum {     SLP_LIFETIME_DEFAULT = 10800,     SLP_LIFETIME_MAXIMUM = 65535 } SLPURLLifetime;</pre> <p>The enumeration <code>SLPURLLifetime</code> contains URL lifetime values, in seconds, that are frequently used. <code>SLP_LIFETIME_DEFAULT</code> is 3 hours, while <code>SLP_LIFETIME_MAXIMUM</code> is 18 hours, which corresponds to the maximum size of the lifetime field in SLP messages. Note that on registration <code>SLP_LIFETIME_MAXIMUM</code> causes the advertisement to be continually reregistered until the process exits.</p> <b>The SLPBoolean Type</b>  <pre>typedef enum {     SLP_FALSE = 0,     SLP_TRUE = 1 } SLPBoolean;</pre> <p>The enumeration <code>SLPBoolean</code> is used as a Boolean flag.</p> <b>The Service URL Structure</b>  <pre>typedef struct srvurl {     char *s_pcSrvType;     char *s_pcHost;     int s_iPort;     char *s_pcNetFamily;     char *s_pcSrvPart; } SLPsrvURL;</pre>	

The `SLPSrvURL` structure is filled in by the `SLPParseSrvURL()` function with information parsed from a character buffer containing a service URL. The fields correspond to different parts of the URL, as follows:

<code>s_pcSrvType</code>	A pointer to a character string containing the service type name, including naming authority.
<code>s_pcHost</code>	A pointer to a character string containing the host identification information.
<code>s_iPort</code>	The port number, or zero, if none. The port is only available if the transport is IP.
<code>s_pcNetFamily</code>	A pointer to a character string containing the network address family identifier. Possible values are "ipx" for the IPX family, "at" for the Appletalk family, and "", the empty string, for the IP address family.
<code>s_pcSrvPart</code>	The remainder of the URL, after the host identification.  The host and port should be sufficient to open a socket to the machine hosting the service; the remainder of the URL should allow further differentiation of the service.

### The SLPHandle

```
typedef void* SLPHandle;
```

The `SLPHandle` type is returned by `SLPOpen()` and is a parameter to all SLP functions. It serves as a handle for all resources allocated on behalf of the process by the SLP library. The type is opaque.

### Callbacks

Include a function pointer to a callback function specific to a particular API operation in the parameter list when the API function is invoked. The callback function is called with the results of the operation in both the synchronous and asynchronous cases. When the callback function is invoked, the memory included in the callback parameters is owned by the API library, and the client code in the callback must copy out the contents if it wants to maintain the information longer than the duration of the current callback call.

Each callback parameter list contains parameters for reporting the results of the operation, as well as an error code parameter and a cookie parameter. The error code parameter reports the error status of the ongoing (for asynchronous) or completed (for synchronous) operation. The cookie parameter allows the client code that starts the operation by invoking the API function to pass information down to the callback without using global variables. The callback returns an `SLPBoolean` to indicate whether the API library should continue processing the operation. If the value returned from the callback is `SLP_TRUE`, asynchronous operations are terminated. Synchronous operations ignore the return since the operation is already complete.

## slp\_api(3SLP)

SLPRegReport ()

```
typedef void SLPRegReport(SLPHandle hSLP,  
    SLPError errCode,  
    void *pvCookie);
```

SLPRegReport () is the callback function to the SLPReg (), SLPDereg (), and SLPDelAttrs () functions. The SLPRegReport () callback has the following parameters:

<i>hSLP</i>	The SLPHandle () used to initiate the operation.
<i>errCode</i>	An error code indicating if an error occurred during the operation.
<i>pvCookie</i>	Memory passed down from the client code that called the original API function, starting the operation. It may be NULL.

SLPSrvTypeCallback ()

```
typedef SLPBoolean SLPSrvTypeCallback(SLPHandle hSLP,  
    const char* pcSrvTypes,  
    SLPError errCode,  
    void *pvCookie);
```

The SLPSrvTypeCallback () type is the type of the callback function parameter to the SLPFindSrvTypes () function. The results are collated when the *hSLP* handle is opened either synchronously or asynchronously. The SLPSrvTypeCallback () callback has the following parameters:

<i>hSLP</i>	The SLPHandle 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 SLP_OK, then the API library may 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. It can be NULL.

**SLPSrvURLCallback**

```
typedef SLPBoolean SLPSrvURLCallback(SLPHandle hSLP,  
    const char* pcSrvURL,  
    unsigned short usLifetime,  
    SLPError errCode,  
    void *pvCookie);
```

The `SLPSrvURLCallback()` type is the type of the callback function parameter to the `SLPFindSrvs()` function. The results are collated, regardless of whether the `hSLP` was opened collated or uncollated. The `SLPSrvURLCallback()` callback has the following parameters:

<i>hSLP</i>	The <code>SLPHandle</code> used to initiate the operation.
<i>pcSrvURL</i>	A character buffer containing the returned service URL.
<i>usLifetime</i>	An unsigned short giving the life time of the service advertisement. The value must be an unsigned integer less than or equal to <code>SLP_LIFETIME_MAXIMUM</code> .
<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 <code>SLP_OK</code> , then the API library may 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. It can be <code>NULL</code> .

#### SLPAttrCallback

```
typedef SLPBoolean SLPAttrCallback(SLPHandle hSLP,
    const char* pcAttrList,
    SLPError errCode,
    void *pvCookie);
```

The `SLPAttrCallback()` type is the type of the callback function parameter to the `SLPFindAttrs()` function.

The behavior of the callback differs depending upon whether the attribute request was by URL or by service type. If the `SLPFindAttrs()` operation was originally called with a URL, the callback is called once, in addition to the last call, regardless of whether the handle was opened asynchronously or synchronously. The `pcAttrList` parameter contains the requested attributes as a comma-separated list. It is empty if no attributes match the original tag list.

If the `SLPFindAttrs()` operation was originally called with a service type, the value of `pcAttrList` and the calling behavior depend upon whether the handle was opened asynchronously or synchronously. If the handle was opened asynchronously, the callback is called every time the API library has results from a remote agent. The `pcAttrList` parameter is collated between calls, and contains a comma-separated list of the results from the agent that immediately returned. If the handle was opened synchronously, the results are collated from all returning agents, the callback is called once, and the `pcAttrList` parameter is set to the collated result.

`SLPAttrCallback()` callback has the following parameters:

<i>hSLP</i>	The <code>SLPHandle</code> used to initiate the operation.
-------------	--

## slp\_api(3SLP)

<i>pcAttrList</i>	A character buffer containing a comma-separated and null-terminated list of attribute id/value assignments, in SLP wire format.
<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 <code>SLP_OK</code> , then the API library may 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. It can be <code>NULL</code> .
<b>ERRORS</b>	An interface that is part of the SLP API may return one of the following values.
<code>SLP_LAST_CALL</code>	The <code>SLP_LAST_CALL</code> code is passed to callback functions when the API library has no more data for them and therefore no further calls will be made to the callback on the currently outstanding operation. The callback uses this to signal the main body of the client code that no more data will be forthcoming on the operation, so that the main body of the client code can break out of data collection loops. On the last call of a callback during both a synchronous and asynchronous call, the error code parameter has value <code>SLP_LAST_CALL</code> , and the other parameters are all <code>NULL</code> . If no results are returned by an API operation, then only one call is made, with the error parameter set to <code>SLP_LAST_CALL</code> .
<code>SLP_OK</code>	The <code>SLP_OK</code> code indicates that the no error occurred during the operation.
<code>SLP_LANGUAGE_NOT_SUPPORTED</code>	No DA or SA has service advertisement information in the language requested, but at least one DA or SA might have information for that service in another language.
<code>SLP_PARSE_ERROR</code>	The SLP message was rejected by a remote SLP agent. The API returns this error only when no information was retrieved, and at least one SA or DA indicated a protocol error. The data supplied through the API may be malformed or damaged in transit.
<code>SLP_INVALID_REGISTRATION</code>	The API may return this error if an attempt to register a service was rejected by all DAs because of a malformed URL or attributes.



SLP_SCOPE_NOT_SUPPORTED	SLP does not return the error if at least one DA accepts the registration.
SLP_AUTHENTICATION_ABSENT	The API returns this error if the UA or SA has been configured with the <code>net.slp.useScopes</code> list of scopes and the SA request did not specify one or more of these allowable scopes, and no others. It may also be returned by a DA if the scope included in a request is not supported by a DA.
SLP_AUTHENTICATION_FAILED	This error arises when the UA or SA failed to send an authenticator for requests or registrations when security is enabled and thus required.
SLP_INVALID_UPDATE	This error arises when a authentication on an SLP message received from a remote SLP agent failed.
SLP_REFRESH_REJECTED	An update for a nonexisting registration was issued, or the update includes a service type or scope different than that in the initial registration.
SLP_NOT_IMPLEMENTED	The SA attempted to refresh a registration more frequently than the minimum refresh interval. The SA should call the appropriate API function to obtain the minimum refresh interval to use.
SLP_BUFFER_OVERFLOW	An outgoing request overflowed the maximum network MTU size. The request should be reduced in size or broken into pieces and tried again.
SLP_NETWORK_TIMED_OUT	An outgoing request overflowed the maximum network MTU size. The request should be reduced in size or broken into pieces and tried again.
SLP_NETWORK_INIT_FAILED	When no reply can be obtained in the time specified by the configured timeout interval, this error is returned.
SLP_MEMORY_ALLOC_FAILED	If the network cannot initialize properly, this error is returned.
	If the API fails to allocate memory, the operation is aborted and returns this.

## slp\_api(3SLP)

### LIST OF ROUTINES

SLP_PARAMETER_BAD	If a parameter passed into an interface is bad, this error is returned.
SLP_NETWORK_ERROR	The failure of networking during normal operations causes this error to be returned.
SLP_INTERNAL_SYSTEM_ERROR	A basic failure of the API causes this error to be returned. This occurs when a system call or library fails. The operation could not recover.
SLP_HANDLE_IN_USE	In the C API, callback functions are not permitted to recursively call into the API on the same <code>SLPHandle</code> , either directly or indirectly. If an attempt is made to do so, this error is returned from the called API function
SLPOpen()	open an SLP handle
SLPClose()	close an open SLP handle
SLPReg()	register a service advertisement
SLPDereg()	deregister a service advertisement
SLPDelAttrs()	delete attributes
SLPFindSrvTypes()	return service types
SLPFindSrvs()	return service URLs
SLPFindAttrs()	return service attributes
SLPGetRefreshInterval()	return the maximum allowed refresh interval for SAs
SLPFindScopes()	return list of configured and discovered scopes
SLPParseSrvURL()	parse service URL
SLPEscape()	escape special characters
SLPUnescape()	translate escaped characters into UTF-8
SLPGetProperty()	return SLP configuration property
SLPSetProperty()	set an SLP configuration property
slp_strerror()	map SLP error code to message
SLPFree()	free memory

### ENVIRONMENT VARIABLES ATTRIBUTES

When `SLP_CONF_FILE` is set, use this file for configuration.  
See `attributes(5)` for descriptions of the following attributes:

slp\_api(3SLP)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWslpu
CSI	CSI-enabled
Interface Stability	Standard
MT-Level	Safe

**SEE ALSO** `slpd(1M)`, `slp.conf(4)`, `slpd.reg(4)`, `attributes(5)`

*System Administration Guide: Resource Management and Network Services*

Guttman, E., Perkins, C., Veizades, J., and Day, M., *RFC 2608, Service Location Protocol, Version 2*, The Internet Society, June 1999.

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

## SLPclose(3SLP)

<b>NAME</b>	SLPclose – close an open SLP handle
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  void <b>SLPclose</b>(SLPHandle <i>phSLP</i>);</pre>
<b>DESCRIPTION</b>	The SLPclose() function frees all resources associated with the handle. If the handle is invalid, the function returns silently. Any outstanding synchronous or asynchronous operations are cancelled, so that their callback functions will not be called any further.
<b>PARAMETERS</b>	<i>phSLP</i> An SLPHandle handle returned from a call to SLPopen().
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPclose()</p> <p>The following example will free all resources associated the handle:</p> <pre>SLPHandle hslp     SLPclose(hslp);</pre>
<b>ENVIRONMENT VARIABLES</b>	SLP_CONF_FILE When set, use this file for configuration.
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWslpu

**SEE ALSO** slpd(1M), slp\_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5)

*System Administration Guide: Resource Management and Network Services*

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

<b>NAME</b>	SLPDelAttrs – delete attributes				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError SLPDelAttrs(SLPHandle hSLP, const char *pcURL, const char     *pcAttrs, SLPRegReport *callback, void *pvCookie);</pre>				
<b>DESCRIPTION</b>	The SLPDelAttrs() function deletes the selected attributes in the locale of the SLPHandle. If no error occurs, the return value is 0. Otherwise, one of the SLPError codes is returned.				
<b>PARAMETERS</b>	<p><i>hSLP</i>                    The language specific SLPHandle to use to delete attributes. It cannot be NULL.</p> <p><i>pcURL</i>                    The URL of the advertisement from which the attributes should be deleted. It cannot be NULL.</p> <p><i>pcAttrs</i>                  A comma-separated list of attribute ids for the attributes to deregister.</p> <p><i>callback</i>                A callback to report the operation's completion status. It cannot be NULL.</p> <p><i>pvCookie</i>                Memory passed to the callback code from the client. It cannot be NULL.</p>				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Deleting Attributes</p> <p>Use the following example to delete the location and dpi attributes for the URL service:printer:lpr://serv/queue1</p> <pre>SLPHandle hSLP; SLPError err; SLPRegReport report;  err = SLPDelAttrs(hSLP, "service:printer:lpr://serv/queue1",     "location,dpi", report, NULL);</pre>				
<b>ENVIRONMENT VARIABLES</b>	SLP_CONF_FILE    When set, use this file for configuration.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWslpu</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	slpd(1M), slp_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5) <i>System Administration Guide: Resource Management and Network Services</i>				

SLPDelAttrs(3SLP)

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

<b>NAME</b>	SLPDereg – deregister the SLP advertisement				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError SLPDereg(SLPHandle hSLP, const char *pcURL, SLPRegReport     callback, void *pvCookie);</pre>				
<b>DESCRIPTION</b>	The SLPDereg() function deregisters the advertisement for URL <i>pcURL</i> in all scopes where the service is registered and in all language locales, not just the locale of the SLPHandle. If no error occurs, the return value is 0. Otherwise, one of the SLPError codes is returned.				
<b>PARAMETERS</b>	<p><i>hSLP</i>                    The language specific SLPHandle to use for deregistering. <i>hSLP</i> cannot be NULL.</p> <p><i>pcURL</i>                    The URL to deregister. The value of <i>pcURL</i> cannot be NULL.</p> <p><i>callback</i>                A callback to report the operation completion status. <i>callback</i> cannot be NULL.</p> <p><i>pvCookie</i>                Memory passed to the callback code from the client. <i>pvCookie</i> can be NULL.</p>				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPDereg()</p> <p>Use the following example to deregister the advertisement for the URL "service:ftp://csserver":</p> <pre>SLPError err; SLPHandle hSLP; SLPRegReport regreport;  err = SLPDereg(hSLP, "service:ftp://csserver", regreport, NULL);</pre>				
<b>ENVIRONMENT VARIABLES</b>	SLP_CONF_FILE    When set, use this file for configuration.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWslpu</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	slpd(1M), slp_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5) <i>System Administration Guide: Resource Management and Network Services</i> Guttman, E., Perkins, C., Veizades, J., and Day, M., <i>RFC 2608, Service Location Protocol, Version 2</i> , The Internet Society, June 1999.				

SLPDereg(3SLP)

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.



<b>NAME</b>	SLPEscape – escapes SLP reserved characters				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError <b>SLPEscape</b>(const char *pInBuf, char** ppcOutBuf, SLPBoolean     isTag);</pre>				
<b>DESCRIPTION</b>	The SLPEscape() function processes the input string in <i>pInBuf</i> and escapes any SLP reserved characters. If the <i>isTag</i> parameter is SLPTrue, it then 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 <i>ppcOutBuf</i> parameter. This buffer should be deallocated using SLPFree(3SLP) when the memory is no longer needed.				
<b>PARAMETERS</b>	<p><i>pInBuf</i>            Pointer to the input buffer to process for escape characters.</p> <p><i>ppcOutBuf</i>        Pointer to a pointer for the output buffer with the SLP reserved characters escaped. It must be freed using SLPFree() when the memory is no longer needed.</p> <p><i>isTag</i>             When true, checks the input buffer for bad tag characters.</p>				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Converting Attribute Tags</p> <p>The following example shows how to convert the attribute tag ,tag-example, to on the wire format:</p> <pre>SLPError err; char* escapedChars;  err = SLPEscape(",tag-example,", &amp;escapedChars, SLP_TRUE);</pre>				
<b>ENVIRONMENT VARIABLES</b>	SLP_CONF_FILE    When set, use this file for configuration.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWslpu</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	<p>slpd(1M), slp_api(3SLP), SLPFree(3SLP), slp.conf(4), slpd.reg(4), attributes(5)</p> <p><i>System Administration Guide: Resource Management and Network Services</i></p> <p>Guttman, E., Perkins, C., Veizades, J., and Day, M., <i>RFC 2608, Service Location Protocol, Version 2</i>, The Internet Society, June 1999.</p>				

SLPEscape(3SLP)

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

<b>NAME</b>	SLPFindAttrs – return service attributes												
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError SLPFindAttrs(SLPHandle hSLP, const char *pcURL, const     char *pcScopeList, const char *pcAttrIds, SLPAtrCallback *callback,     void *pvCookie);</pre>												
<b>DESCRIPTION</b>	<p>The SLPFindAttrs () function returns service attributes matching the attribute tags for the indicated full or partial URL. If <i>pcURL</i> is a complete URL, the attribute information returned is for that particular service in the language locale of the SLPHandle. If <i>pcURL</i> is a service type, then all attributes for the service type are returned, regardless of the language of registration. Results are returned through the <i>callback</i> parameter.</p> <p>The result is filtered with an SLP attribute request filter string parameter, the syntax of which is described in <i>RFC 2608</i>. If the filter string is the empty string, "", all attributes are returned.</p> <p>If an error occurs in starting the operation, one of the SLPError codes is returned.</p>												
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="vertical-align: top;"><i>hSLP</i></td> <td>The language-specific SLPHandle on which to search for attributes. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>pcURL</i></td> <td>The full or partial URL. See <i>RFC 2608</i> for partial URL syntax. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>pcScopeList</i></td> <td>A pointer to a char containing a comma-separated list of scope names. It cannot be NULL or an empty string, "".</td> </tr> <tr> <td style="vertical-align: top;"><i>pcAttrIds</i></td> <td>The filter string indicating which attribute values to return. Use empty string "" to indicate all values. Wildcards matching all attribute ids having a particular prefix or suffix are also possible. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>callback</i></td> <td>A callback function through which the results of the operation are reported. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>pvCookie</i></td> <td>Memory passed to the callback code from the client. It may be NULL.</td> </tr> </table>	<i>hSLP</i>	The language-specific SLPHandle on which to search for attributes. It cannot be NULL.	<i>pcURL</i>	The full or partial URL. See <i>RFC 2608</i> for partial URL syntax. It cannot be NULL.	<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names. It cannot be NULL or an empty string, "".	<i>pcAttrIds</i>	The filter string indicating which attribute values to return. Use empty string "" to indicate all values. Wildcards matching all attribute ids having a particular prefix or suffix are also possible. It cannot be NULL.	<i>callback</i>	A callback function through which the results of the operation are reported. It cannot be NULL.	<i>pvCookie</i>	Memory passed to the callback code from the client. It may be NULL.
<i>hSLP</i>	The language-specific SLPHandle on which to search for attributes. It cannot be NULL.												
<i>pcURL</i>	The full or partial URL. See <i>RFC 2608</i> for partial URL syntax. It cannot be NULL.												
<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names. It cannot be NULL or an empty string, "".												
<i>pcAttrIds</i>	The filter string indicating which attribute values to return. Use empty string "" to indicate all values. Wildcards matching all attribute ids having a particular prefix or suffix are also possible. It cannot be NULL.												
<i>callback</i>	A callback function through which the results of the operation are reported. It cannot be NULL.												
<i>pvCookie</i>	Memory passed to the callback code from the client. It may be NULL.												
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in <i>slp_api(3SLP)</i> .												
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Returning Service Attributes for a Specific URL</p> <p>Use the following example to return the attributes "location" and "dpi" for the URL "service:printer:lpr://serv/queue1" through the callback <i>attrReturn</i>:</p> <pre>SLPHandle hSLP; SLPAtrCallback attrReturn; SLPError err;</pre>												

## SLPFindAttrs(3SLP)

**EXAMPLE 1** Returning Service Attributes for a Specific URL *(Continued)*

```
err = SLPFindAttrs(hSLP "service:printer:lpr://serv/queue1",
                  "default", "location,dpi", attrReturn, err);
```

**EXAMPLE 2** Returning Service Attributes for All URLs of a Specific Type

Use the following example to return the attributes "location" and "dpi" for all service URLs having type "service:printer:lpr":

```
err = SLPFindAttrs(hSLP, "service:printer:lpr",
                  "default", "location, pi",
                  attrReturn, NULL);
```

### ENVIRONMENT VARIABLES ATTRIBUTES

SLP\_CONF\_FILE When set, use this file for configuration.

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWslpu

**SEE ALSO** [slpd\(1M\)](#), [slp\\_api\(3SLP\)](#), [slp.conf\(4\)](#), [slpd.reg\(4\)](#), [attributes\(5\)](#)

*System Administration Guide: Resource Management and Network Services*

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

<b>NAME</b>	SLPFindScopes – return list of configured and discovered scopes				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError SLPFindScopes (SLPHandle hSLP, char** ppcScopes);</pre>				
<b>DESCRIPTION</b>	<p>The SLPFindScopes() function sets the <i>ppcScopes</i> parameter to a pointer to a comma-separated list including all available scope names. The list of scopes comes from a variety of sources: the configuration file, the <code>net.slp.useScopes</code> property and the <code>net.slp.DAAddresses</code> property, DHCP, or through the DA discovery process. If there is any order to the scopes, preferred scopes are listed before less desirable scopes. There is always at least one string in the array, the default scope, <code>DEFAULT</code>.</p> <p>If no error occurs, SLPFindScopes() returns <code>SLP_OK</code>, otherwise, it returns the appropriate error code.</p>				
<b>PARAMETERS</b>	<p><i>hSLP</i>                    The SLPHandle on which to search for scopes. <i>hSLP</i> cannot be NULL.</p> <p><i>ppcScopes</i>                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(). See SLPFree(3SLP)</p>				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in <code>slp_api(3SLP)</code> .				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Finding Configured or Discovered Scopes</p> <p>Use the following example to find configured or discovered scopes:</p> <pre>SLPHandle hSLP; char *ppcScopes; SLPError err;  error = SLPFindScopes(hSLP, &amp; ppcScopes);</pre>				
<b>ENVIRONMENT VARIABLES</b>	<code>SLP_CONF_FILE</code> When set, use this file for configuration.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWslpu</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	<p><code>slpd(1M)</code>, <code>slp_api(3SLP)</code>, <code>SLPFree(3SLP)</code>, <code>slp.conf(4)</code>, <code>slpd.reg(4)</code>, <code>attributes(5)</code></p> <p><i>System Administration Guide: Resource Management and Network Services</i></p> <p>Guttman, E., Perkins, C., Veizades, J., and Day, M., <i>RFC 2608, Service Location Protocol, Version 2</i>, The Internet Society, June 1999.</p>				

SLPFindScopes(3SLP)

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

<b>NAME</b>	SLPFindSrvs – return service URLs												
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError <b>SLPFindSrvs</b>(SLPHandle <i>hSLP</i>, const char *<i>pcServiceType</i>, const char *<i>pcScopeList</i>, const char *<i>pcSearchFilter</i>, SLPSrvURLCallback *<i>callback</i>, void *<i>pvCookie</i>);</pre>												
<b>DESCRIPTION</b>	<p>The SLPFindSrvs () function issues a request for SLP services. The query is for services on a language-specific SLPHandle. It returns the results through the <i>callback</i>. The parameters will determine the results.</p> <p>If an error occurs in starting the operation, one of the SLPError codes is returned.</p>												
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="vertical-align: top;"><i>hSLP</i></td> <td>The language-specific SLPHandle on which to search for services. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>pcServiceType</i></td> <td> <p>The service type string for the request. The <i>pcServiceType</i> can be discovered by a call to SLPSrvTypes (). Examples of service type strings include</p> <p>"service:printer:lpr"</p> <p>or</p> <p>"service:nfs"</p> <p><i>pcServiceType</i> cannot be NULL.</p> </td> </tr> <tr> <td style="vertical-align: top;"><i>pcScopeList</i></td> <td>A pointer to a char containing a comma-separated list of scope names. It cannot be NULL or an empty string, "".</td> </tr> <tr> <td style="vertical-align: top;"><i>pcSearchFilter</i></td> <td>A query formulated of attribute pattern matching expressions in the form of a LDAPv3 search filter. See RFC 2254. If this filter is empty, "", all services of the requested type in the specified scopes are returned. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>callback</i></td> <td>A callback through which the results of the operation are reported. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>pvCookie</i></td> <td>Memory passed to the callback code from the client. It can be NULL.</td> </tr> </table>	<i>hSLP</i>	The language-specific SLPHandle on which to search for services. It cannot be NULL.	<i>pcServiceType</i>	<p>The service type string for the request. The <i>pcServiceType</i> can be discovered by a call to SLPSrvTypes (). Examples of service type strings include</p> <p>"service:printer:lpr"</p> <p>or</p> <p>"service:nfs"</p> <p><i>pcServiceType</i> cannot be NULL.</p>	<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names. It cannot be NULL or an empty string, "".	<i>pcSearchFilter</i>	A query formulated of attribute pattern matching expressions in the form of a LDAPv3 search filter. See RFC 2254. If this filter is empty, "", all services of the requested type in the specified scopes are returned. It cannot be NULL.	<i>callback</i>	A callback through which the results of the operation are reported. It cannot be NULL.	<i>pvCookie</i>	Memory passed to the callback code from the client. It can be NULL.
<i>hSLP</i>	The language-specific SLPHandle on which to search for services. It cannot be NULL.												
<i>pcServiceType</i>	<p>The service type string for the request. The <i>pcServiceType</i> can be discovered by a call to SLPSrvTypes (). Examples of service type strings include</p> <p>"service:printer:lpr"</p> <p>or</p> <p>"service:nfs"</p> <p><i>pcServiceType</i> cannot be NULL.</p>												
<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names. It cannot be NULL or an empty string, "".												
<i>pcSearchFilter</i>	A query formulated of attribute pattern matching expressions in the form of a LDAPv3 search filter. See RFC 2254. If this filter is empty, "", all services of the requested type in the specified scopes are returned. It cannot be NULL.												
<i>callback</i>	A callback through which the results of the operation are reported. It cannot be NULL.												
<i>pvCookie</i>	Memory passed to the callback code from the client. It can be NULL.												
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).												
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPFindSrvs ()</p> <p>The following example finds all advertisements for printers supporting the LPR protocol with the dpi attribute 300 in the default scope:</p>												

## SLPFindSrvs(3SLP)

**EXAMPLE 1** Using SLPFindSrvs () (Continued)

```
SLPError err;
SLPHandle hSLP;
SLPSrvURLCallback srvngst;

err = SLPFindSrvs(hSLP,
                 "service:printer:lpr",
                 "default",
                 "(dpi=300)",
                 srvngst,
                 NULL);
```

### ENVIRONMENT VARIABLES ATTRIBUTES

SLP\_CONF\_FILE When set, use this file for configuration.

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWslpu

### SEE ALSO

slpd(1M), slp\_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5)

*System Administration Guide: Resource Management and Network Services*

Howes, T. *RFC 2254, The String Representation of LDAP Search Filters*. The Internet Society. 1997.

Guttman, E., Perkins, C., Veizades, J., and Day, M. *RFC 2608, Service Location Protocol, Version 2*. The Internet Society. June 1999.

Kempf, J. and Guttman, E. *RFC 2614, An API for Service Location*. The Internet Society. June 1999.



<b>NAME</b>	SLPFindSrvTypes – find service types										
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError SLPFindSrvTypes (SLPHandle hSLP, const char     *pcNamingAuthority, const char *pcScopeList, SLPsrvTypeCallback     *callback, void *pvCookie);</pre>										
<b>DESCRIPTION</b>	<p>The SLPFindSrvTypes() function 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 indicated naming authority.</p> <p>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; the SLP_PARAMETER_BAD error code will be returned if you include it explicitly.</p> <p>The service type names are returned with the naming authority included in the following format:</p> <pre>service-type "." naming-authority</pre> <p>unless the naming authority is the default, in which case, just the service type name is returned.</p> <p>If an error occurs in starting the operation, one of the SLPError codes is returned.</p>										
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="vertical-align: top;"><i>hSLP</i></td> <td>The SLPHandle on which to search for types. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>pcNamingAuthority</i></td> <td>The naming authority to search. Use "*" to search all naming authorities; use the empty string "" to search the default naming authority. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>pcScopeList</i></td> <td>A pointer to a char containing a comma-separated list of scope names to search for service types. It cannot be NULL or an empty string, "".</td> </tr> <tr> <td style="vertical-align: top;"><i>callback</i></td> <td>A callback through which the results of the operation are reported. It cannot be NULL.</td> </tr> <tr> <td style="vertical-align: top;"><i>pvCookie</i></td> <td>Memory passed to the callback code from the client. It can be NULL.</td> </tr> </table>	<i>hSLP</i>	The SLPHandle on which to search for types. It cannot be NULL.	<i>pcNamingAuthority</i>	The naming authority to search. Use "*" to search all naming authorities; use the empty string "" to search the default naming authority. It cannot be NULL.	<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names to search for service types. It cannot be NULL or an empty string, "".	<i>callback</i>	A callback through which the results of the operation are reported. It cannot be NULL.	<i>pvCookie</i>	Memory passed to the callback code from the client. It can be NULL.
<i>hSLP</i>	The SLPHandle on which to search for types. It cannot be NULL.										
<i>pcNamingAuthority</i>	The naming authority to search. Use "*" to search all naming authorities; use the empty string "" to search the default naming authority. It cannot be NULL.										
<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names to search for service types. It cannot be NULL or an empty string, "".										
<i>callback</i>	A callback through which the results of the operation are reported. It cannot be NULL.										
<i>pvCookie</i>	Memory passed to the callback code from the client. It can be NULL.										
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).										
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPFindSrvTypes()</p> <p>The following example finds all service type names in the default scope and default naming authority:</p> <pre>SLPError err; SLPHandle hSLP;</pre>										

## SLPFindSrvTypes(3SLP)

**EXAMPLE 1** Using SLPFindSrvTypes () (Continued)

```
SLPSrvTypeCallback findsrvtypes;  
err = SLPFindSrvTypes(hSLP, "", "default", findsrvtypes, NULL);
```

### ENVIRONMENT VARIABLES ATTRIBUTES

SLP\_CONF\_FILE When set, use this file for configuration.

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWslpu

### SEE ALSO

slpd(1M), slp\_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5)

*System Administration Guide: Resource Management and Network Services*

Guttman, E., Perkins, C., Veizades, J., and Day, M., *RFC 2608, Service Location Protocol, Version 2*, The Internet Society, June 1999.

Howes, T., *RFC 2254, The String Representation of LDAP Search Filters*, The Internet Society, 1997.

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

<b>NAME</b>	SLPFree – frees memory				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError <b>SLPFree</b>(void *pvMem) ;</pre>				
<b>DESCRIPTION</b>	The SLPFree() function frees memory returned from SLPParseSrvURL(), SLPFindScopes(), SLPEscape(), and SLPUnescape().				
<b>PARAMETERS</b>	<i>pvMem</i> A pointer to the storage allocated by the SLPParseSrvURL(), SLPFindScopes(), SLPEscape(), and SLPUnescape() functions. <i>pvMem</i> is ignored if its value is NULL.				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPFree()</p> <p>The following example illustrates how to call SLPFree(). It assumes that SrvURL contains previously allocated memory.</p> <pre>SLPError err;  err = SLPFree((void*) SrvURL);</pre>				
<b>ENVIRONMENT VARIABLES</b>	SLP_CONF_FILE    When set, use this file for configuration.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWslpu</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	<p>slpd(1M), SLPEscape(3SLP), SLPFindScopes(3SLP), SLPParseSrvURL(3SLP), SLPUnescape(3SLP), slp_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5)</p> <p><i>System Administration Guide: Resource Management and Network Services</i></p> <p>Guttman, E., Perkins, C., Veizades, J., and Day, M., <i>RFC 2608, Service Location Protocol, Version 2</i>, The Internet Society, June 1999.</p> <p>Kempf, J. and Guttman, E., <i>RFC 2614, An API for Service Location</i>, The Internet Society, June 1999.</p>				

## SLPGetProperty(3SLP)

<b>NAME</b>	SLPGetProperty – return SLP configuration property				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  const char* <b>SLPGetProperty</b>(const char* <i>pcName</i>);</pre>				
<b>DESCRIPTION</b>	The SLPGetProperty() function returns the value of the corresponding SLP property name, or NULL, if none. If there is no error, SLPGetProperty() returns a pointer to the property value. If the property was not set, it returns the empty string, "". If an error occurs, SLPGetProperty() returns NULL. The returned string should not be freed.				
<b>PARAMETERS</b>	<i>pcName</i> A null-terminated string with the property name. <i>pcName</i> cannot be NULL.				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPGetProperty()</p> <p>Use the following example to return a list of configured scopes:</p> <pre>const char* useScopes  useScopes = SLPGetProperty("net.slp.useScopes");</pre>				
<b>ENVIRONMENT VARIABLES</b>	SLP_CONF_FILE When set, use this file for configuration.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWslpu</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	slpd(1M), slp_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5) <i>System Administration Guide: Resource Management and Network Services</i> Kempf, J. and Guttman, E., <i>RFC 2614, An API for Service Location</i> , The Internet Society, June 1999.				

<b>NAME</b>	SLPGetRefreshInterval – return the maximum allowed refresh interval				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  int <b>SLPGetRefreshInterval</b> (void);</pre>				
<b>DESCRIPTION</b>	The SLPGetRefreshInterval () function returns the maximum across all DAs of the min-refresh-interval attribute. This value satisfies the advertised refresh interval bounds for all DAs. If this value is used by the SA, it assures that no refresh registration will be rejected. If no DA advertises a min-refresh-interval attribute, a value of 0 is returned. If an error occurs, an SLP error code is returned.				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPGetRefreshInterval ()</p> <p>Use the following example to return the maximum valid refresh interval for SA:</p> <pre>int minrefresh  minrefresh = SLPGetRefreshInterval( );</pre>				
<b>ENVIRONMENT VARIABLES</b>	SLP_CONF_FILE When set, use this file for configuration.				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWslpu</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	slpd(1M), slp_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5) <i>System Administration Guide: Resource Management and Network Services</i> Kempf, J. and Guttman, E., <i>RFC 2614, An API for Service Location</i> , The Internet Society, June 1999.				

## SLPOpen(3SLP)

<b>NAME</b>	SLPOpen – open an SLP handle						
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError <b>SLPOpen</b>(const char *<i>pcLang</i>, SLPBoolean <i>isAsync</i>, SLPHandle                   *<i>phSLP</i>);</pre>						
<b>DESCRIPTION</b>	<p>The SLPOpen() function returns a SLPHandle handle in the <i>phSLP</i> parameter for the language locale passed in as the <i>pcLang</i> parameter. The client indicates if operations on the handle are to be synchronous or asynchronous through the <i>isAsync</i> parameter. The handle encapsulates the language locale for SLP requests issued through the handle, and any other resources required by the implementation. 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 <i>phSLP</i> parameter is NULL.</p> <p>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() function terminates any outstanding calls on the handle.</p>						
<b>PARAMETERS</b>	<table><tr><td><i>pcLang</i></td><td>A pointer to an array of characters containing the language tag set forth in RFC 1766 for the natural language locale of requests issued on the handle. This parameter cannot be NULL.</td></tr><tr><td><i>isAsync</i></td><td>An SLPBoolean indicating whether or not the SLPHandle should be opened for an asynchronous operation.</td></tr><tr><td><i>phSLP</i></td><td>A pointer to an SLPHandle in which the open SLPHandle is returned. If an error occurs, the value upon return is NULL.</td></tr></table>	<i>pcLang</i>	A pointer to an array of characters containing the language tag set forth in RFC 1766 for the natural language locale of requests issued on the handle. This parameter cannot be NULL.	<i>isAsync</i>	An SLPBoolean indicating whether or not the SLPHandle should be opened for an asynchronous operation.	<i>phSLP</i>	A pointer to an SLPHandle in which the open SLPHandle is returned. If an error occurs, the value upon return is NULL.
<i>pcLang</i>	A pointer to an array of characters containing the language tag set forth in RFC 1766 for the natural language locale of requests issued on the handle. This parameter cannot be NULL.						
<i>isAsync</i>	An SLPBoolean indicating whether or not the SLPHandle should be opened for an asynchronous operation.						
<i>phSLP</i>	A pointer to an SLPHandle in which the open SLPHandle is returned. If an error occurs, the value upon return is NULL.						
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).						
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPOpen()</p> <p>Use the following example to open a synchronous handle for the German (“de”) locale:</p> <pre>SLPHandle HSLP; SLPError err;  err = SLPOpen("de", SLP_FALSE, &amp;HSLP)</pre>						
<b>ENVIRONMENT VARIABLES</b>	SLP_CONF_FILE When set, use this file for configuration.						
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:						

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWslpu

**SEE ALSO** `slpd(1M)`, `slp_api(3SLP)`, `slp.conf(4)`, `slpd.reg(4)`, `attributes(5)`

*System Administration Guide: Resource Management and Network Services*

Alvestrand, H., *RFC 1766, Tags for the Identification of Languages*, Network Working Group, March 1995.

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

## SLPParseSrvURL(3SLP)

<b>NAME</b>	SLPParseSrvURL – parse service URL				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError <b>SLPParseSrvURL</b>(const char *pcSrvURL, SLPSrvURL**     ppSrvURL);</pre>				
<b>DESCRIPTION</b>	<p>The SLPParseSrvURL() routine parses the URL passed in as the argument into a service URL structure and returns it in the ppSrvURL pointer. If a parser error occurs, returns SLP_PARSE_ERROR. The structure returned in ppSrvURL should be freed with SLPFree(). If the URL has no service part, the s_pcSrvPart string is the empty string, "", that is, it is 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_pcNetFamily field is the empty string.</p> <p>If no error occurs, the return value is the SLP_OK. Otherwise, if an error occurs, one of the SLPError codes is returned.</p>				
<b>PARAMETERS</b>	<p><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. It may not be NULL.</p> <p><i>ppSrvURL</i>            A pointer to a pointer for the SLPSrvURL structure to receive the parsed URL. It may not be NULL.</p>				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPParseSrvURL()</p> <p>The following example uses the SLPParseSrvURL() function to parse the service URL service:printer:lpr://serv/queue1:</p> <pre>SLPSrvURL* surl; SLPError err;  err = SLPParseSrvURL("service:printer:lpr://serv/queue1", &amp;surl);</pre>				
<b>ENVIRONMENT VARIABLES ATTRIBUTES</b>	<p>SLP_CONF_FILE    When set, use this file for configuration.</p> <p>See attributes(5) for descriptions of the following attributes:</p>				
<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Availability</td> <td style="text-align: center;">SUNWslpu</td> </tr> </tbody> </table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	<p>slpd(1M), slp_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5)</p> <p><i>System Administration Guide: Resource Management and Network Services</i></p>				



SLPParseSrvURL(3SLP)

Guttman, E., Perkins, C., Veizades, J., and Day, M., *RFC 2608, Service Location Protocol, Version 2*, The Internet Society, June 1999.

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

## SLPReg(3SLP)

<b>NAME</b>	SLPReg – register an SLP advertisement																
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError <b>SLPReg</b>(SLPHandle <i>hSLP</i>, const char *<i>pcSrvURL</i>, const     unsigned short <i>usLifetime</i>, const char *<i>pcSrvType</i>, const char     *<i>pcAttrs</i>, SLPBoolean <i>fresh</i>, SLPRegReport <i>callback</i>, void *<i>pvCookie</i>);</pre>																
<b>DESCRIPTION</b>	<p>The SLPReg() function registers the URL in <i>pcSrvURL</i> having the lifetime <i>usLifetime</i> with the attribute list in <i>pcAttrs</i>. The <i>pcAttrs</i> list is a comma-separated list of attribute assignments in on-the-wire format (including escaping of reserved characters). The <i>usLifetime</i> parameter must be nonzero and less than or equal to SLP_LIFETIME_MAXIMUM. If the fresh flag is SLP_TRUE, then the registration is new, the SLP protocol <i>fresh</i> flag is set, and the registration replaces any existing registrations.</p> <p>The <i>pcSrvType</i> parameter is a service type name and can be included for service URLs that are not in the service: scheme. If the URL is in the service: scheme, the <i>pcSrvType</i> parameter is ignored. If the fresh flag is SLP_FALSE, then an existing registration is updated. Rules for new and updated registrations, and the format for <i>pcAttrs</i> and <i>pcScopeList</i>, can be found in RFC 2608. Registrations and updates take place in the language locale of the <i>hSLP</i> handle.</p> <p>The API library is required to perform the operation in all scopes obtained through configuration.</p>																
<b>PARAMETERS</b>	<table><tr><td><i>hSLP</i></td><td>The language specific SLPHandle on which to register the advertisement. <i>hSLP</i> cannot be NULL.</td></tr><tr><td><i>pcSrvURL</i></td><td>The URL to register. The value of <i>pcSrvURL</i> cannot be NULL or the empty string.</td></tr><tr><td><i>usLifetime</i></td><td>An unsigned short giving the life time of the service advertisement, in seconds. The value must be an unsigned integer less than or equal to SLP_LIFETIME_MAXIMUM.</td></tr><tr><td><i>pcSrvType</i></td><td>The service type. If <i>pcURL</i> is a service: URL, then this parameter is ignored. <i>pcSrvType</i> cannot be NULL.</td></tr><tr><td><i>pcAttrs</i></td><td>A comma-separated list of attribute assignment expressions for the attributes of the advertisement. <i>pcAttrs</i> cannot be NULL. Use the empty string, "", to indicate no attributes.</td></tr><tr><td><i>fresh</i></td><td>An SLPBoolean that is SLP_TRUE if the registration is new or SLP_FALSE if it is a reregistration.</td></tr><tr><td><i>callback</i></td><td>A callback to report the operation completion status. <i>callback</i> cannot be NULL.</td></tr><tr><td><i>pvCookie</i></td><td>Memory passed to the callback code from the client. <i>pvCookie</i> can be NULL.</td></tr></table>	<i>hSLP</i>	The language specific SLPHandle on which to register the advertisement. <i>hSLP</i> cannot be NULL.	<i>pcSrvURL</i>	The URL to register. The value of <i>pcSrvURL</i> cannot be NULL or the empty string.	<i>usLifetime</i>	An unsigned short giving the life time of the service advertisement, in seconds. The value must be an unsigned integer less than or equal to SLP_LIFETIME_MAXIMUM.	<i>pcSrvType</i>	The service type. If <i>pcURL</i> is a service: URL, then this parameter is ignored. <i>pcSrvType</i> cannot be NULL.	<i>pcAttrs</i>	A comma-separated list of attribute assignment expressions for the attributes of the advertisement. <i>pcAttrs</i> cannot be NULL. Use the empty string, "", to indicate no attributes.	<i>fresh</i>	An SLPBoolean that is SLP_TRUE if the registration is new or SLP_FALSE if it is a reregistration.	<i>callback</i>	A callback to report the operation completion status. <i>callback</i> cannot be NULL.	<i>pvCookie</i>	Memory passed to the callback code from the client. <i>pvCookie</i> can be NULL.
<i>hSLP</i>	The language specific SLPHandle on which to register the advertisement. <i>hSLP</i> cannot be NULL.																
<i>pcSrvURL</i>	The URL to register. The value of <i>pcSrvURL</i> cannot be NULL or the empty string.																
<i>usLifetime</i>	An unsigned short giving the life time of the service advertisement, in seconds. The value must be an unsigned integer less than or equal to SLP_LIFETIME_MAXIMUM.																
<i>pcSrvType</i>	The service type. If <i>pcURL</i> is a service: URL, then this parameter is ignored. <i>pcSrvType</i> cannot be NULL.																
<i>pcAttrs</i>	A comma-separated list of attribute assignment expressions for the attributes of the advertisement. <i>pcAttrs</i> cannot be NULL. Use the empty string, "", to indicate no attributes.																
<i>fresh</i>	An SLPBoolean that is SLP_TRUE if the registration is new or SLP_FALSE if it is a reregistration.																
<i>callback</i>	A callback to report the operation completion status. <i>callback</i> cannot be NULL.																
<i>pvCookie</i>	Memory passed to the callback code from the client. <i>pvCookie</i> can be NULL.																

**ERRORS** This function or its callback may return any SLP error code. See the **ERRORS** section in `slp_api(3SLP)`.

**EXAMPLES** **EXAMPLE 1** An Initial Registration

The following example shows an initial registration for the "service:video://bldg15" camera service for three hours:

```
SLPError err;
SLPHandle hSLP;
SLPRegReport regreport;
err = SLPReg(hSLP, "service:video://bldg15",
            10800, "", "(location=B15-corridor)",
            (scan-rate=100)", SLP_TRUE,
            regRpt, NULL);
```

**ENVIRONMENT VARIABLES** `SLP_CONF_FILE` When set, use this file for configuration.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWslpu

**SEE ALSO** `slpd(1M)`, `slp_api(3SLP)`, `slp.conf(4)`, `slpd.reg(4)`, `attributes(5)`

*System Administration Guide: Resource Management and Network Services*

Guttman, E., Perkins, C., Veizades, J., and Day, M., *RFC 2608, Service Location Protocol, Version 2*, The Internet Society, June 1999.

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

## SLPsetProperty(3SLP)

<b>NAME</b>	SLPsetProperty – set an SLP configuration property				
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  void <b>SLPsetProperty</b>(const char *pcName, const char *pcValue);</pre>				
<b>DESCRIPTION</b>	The SLPsetProperty() function sets the value of the SLP property to the new value. The pcValue parameter contains the property value as a string.				
<b>PARAMETERS</b>	<table><tr><td>pcName</td><td>A null-terminated string with the property name. pcName cannot be NULL.</td></tr><tr><td>pcValue</td><td>A null-terminated string with the property value. pcValue cannot be NULL.</td></tr></table>	pcName	A null-terminated string with the property name. pcName cannot be NULL.	pcValue	A null-terminated string with the property value. pcValue cannot be NULL.
pcName	A null-terminated string with the property name. pcName cannot be NULL.				
pcValue	A null-terminated string with the property value. pcValue cannot be NULL.				
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in slp_api(3SLP).				
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Setting a Configuration Property</p> <p>The following example shows to set the property net.slp.typeHint to service:ftp:</p> <pre>SLPsetProperty ("net.slp.typeHint" "service:ftp");</pre>				
<b>ENVIRONMENT VARIABLES ATTRIBUTES</b>	<p>SLP_CONF_FILE When set, use this file for configuration.</p> <p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWslpu</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWslpu				
<b>SEE ALSO</b>	slpd(1M), slp_api(3SLP), slp.conf(4), slpd.reg(4), attributes(5) <i>System Administration Guide: Resource Management and Network Services</i> Kempf, J. and Guttman, E., <i>RFC 2614, An API for Service Location</i> , The Internet Society, June 1999.				

<b>NAME</b>	slp_strerror – map SLP error codes to messages
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  const char* <b>slp_strerror</b>(SLPError <i>err_code</i>);</pre>
<b>DESCRIPTION</b>	The <code>slp_strerror()</code> function maps <code>err_code</code> to a string explanation of the error. The returned string is owned by the library and must not be freed.
<b>PARAMETERS</b>	<i>err_code</i> An SLP error code.
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in <code>slp_api(3SLP)</code> .
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using <code>slp_strerror()</code></p> <p>The following example returns the message that corresponds to the error code:</p> <pre>SLPError error; const char* msg; msg = slp_strerror(err);</pre>
<b>ENVIRONMENT VARIABLES</b>	<code>SLP_CONF_FILE</code> When set, use this file for configuration.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWslpu

**SEE ALSO** `slpd(1M)`, `slp_api(3SLP)`, `slp.conf(4)`, `slpd.reg(4)`, `attributes(5)`

*System Administration Guide: Resource Management and Network Services*

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

## SLPUnescape(3SLP)

<b>NAME</b>	SLPUnescape – translate escaped characters into UTF-8						
<b>SYNOPSIS</b>	<pre>#include &lt;slp.h&gt;  SLPError <b>SLPUnescape</b>(const char *pcInBuf, char** ppcOutBuf,                       SLPBoolean isTag);</pre>						
<b>DESCRIPTION</b>	The SLPUnescape () function processes the input string in <i>pcInBuf</i> and unescapes any SLP reserved characters. If the <i>isTag</i> parameter is <i>SLPTrue</i> , then look for bad tag characters and signal an error if any are found with the <i>SLP_PARSE_ERROR</i> code. No transformation is performed if the input string is an opaque. The results are put into a buffer allocated by the API library and returned in the <i>ppcOutBuf</i> parameter. This buffer should be deallocated using <i>SLPFree(3SLP)</i> when the memory is no longer needed.						
<b>PARAMETERS</b>	<table border="0"> <tr> <td style="padding-right: 20px;"><i>pcInBuf</i></td> <td>Pointer to the input buffer to process for escape characters.</td> </tr> <tr> <td><i>ppcOutBuf</i></td> <td>Pointer to a pointer for the output buffer with the SLP reserved characters escaped. Must be freed using <i>SLPFree(3SLP)</i> when the memory is no longer needed.</td> </tr> <tr> <td><i>isTag</i></td> <td>When true, the input buffer is checked for bad tag characters.</td> </tr> </table>	<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 SLP reserved characters escaped. Must be freed using <i>SLPFree(3SLP)</i> when the memory is no longer needed.	<i>isTag</i>	When true, the input buffer is checked for bad tag characters.
<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 SLP reserved characters escaped. Must be freed using <i>SLPFree(3SLP)</i> when the memory is no longer needed.						
<i>isTag</i>	When true, the input buffer is checked for bad tag characters.						
<b>ERRORS</b>	This function or its callback may return any SLP error code. See the ERRORS section in <i>slp_api(3SLP)</i> .						
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using SLPUnescape ()</p> <p>The following example decodes the representation for ", tag, ":</p> <pre>char* pcOutBuf; SLPError err;  err = SLPUnescape("\\2c tag\\2c", &amp;pcOutbuf, SLP_TRUE);</pre>						
<b>ENVIRONMENT VARIABLES ATTRIBUTES</b>	<p><i>SLP_CONF_FILE</i> When set, use this file for configuration.</p> <p>See <i>attributes(5)</i> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWslpu</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWslpu		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWslpu						
<b>SEE ALSO</b>	<p><i>slpd(1M)</i>, <i>SLPFree(3SLP)</i>, <i>slp_api(3SLP)</i>, <i>slp.conf(4)</i>, <i>slpd.reg(4)</i>, <i>attributes(5)</i></p> <p><i>System Administration Guide: Resource Management and Network Services</i></p> <p>Guttman, E., Perkins, C., Veizades, J., and Day, M., <i>RFC 2608, Service Location Protocol, Version 2</i>, The Internet Society, June 1999.</p>						

SLPUnescape(3SLP)

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

## socket(3SOCKET)

<b>NAME</b>	socket – create an endpoint for communication								
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lsocket -lnsl [ <i>library</i> ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  int <b>socket</b>(int <i>domain</i>, int <i>type</i>, int <i>protocol</i>);</pre>								
<b>DESCRIPTION</b>	<p>socket () creates an endpoint for communication and returns a descriptor.</p> <p>The <i>domain</i> parameter specifies a communications domain within which communication will take place; this selects the protocol family which should be used. The protocol family generally is the same as the address family for the addresses supplied in later operations on the socket. These families are defined in the include file &lt;sys/socket.h&gt;. There must be an entry in the netconfig(4) file for at least each protocol family and type required. If <i>protocol</i> has been specified, but no exact match for the tuple family, type, protocol is found, then the first entry containing the specified family and type with zero for protocol will be used. The currently understood formats are:</p> <table><tr><td>PF_UNIX</td><td>UNIX system internal protocols</td></tr><tr><td>PF_INET</td><td>Internet Protocol Version 4 (IPv4)</td></tr><tr><td>PF_INET6</td><td>Internet Protocol Version 6 (IPv6)</td></tr><tr><td>PF_NCA</td><td>Network Cache and Accelerator (NCA) protocols</td></tr></table> <p>The socket has the indicated <i>type</i>, which specifies the communication semantics. Currently defined types are:</p> <pre>SOCK_STREAM SOCK_DGRAM SOCK_RAW SOCK_SEQPACKET SOCK_RDM</pre> <p>A SOCK_STREAM type provides sequenced, reliable, two-way connection-based byte streams. An out-of-band data transmission mechanism may be supported. A SOCK_DGRAM socket supports datagrams (connectionless, unreliable messages of a fixed (typically small) maximum length). A SOCK_SEQPACKET socket may provide a sequenced, reliable, two-way connection-based data transmission path for datagrams of fixed maximum length; a consumer may be required to read an entire packet with each read system call. This facility is protocol specific, and presently not implemented for any protocol family. SOCK_RAW sockets provide access to internal network interfaces. The types SOCK_RAW, which is available only to the superuser, and SOCK_RDM, for which no implementation currently exists, are not described here.</p> <p><i>protocol</i> specifies a particular protocol to be used with the socket. Normally only a single protocol exists to support a particular socket type within a given protocol family. However, multiple protocols may exist, in which case a particular protocol</p>	PF_UNIX	UNIX system internal protocols	PF_INET	Internet Protocol Version 4 (IPv4)	PF_INET6	Internet Protocol Version 6 (IPv6)	PF_NCA	Network Cache and Accelerator (NCA) protocols
PF_UNIX	UNIX system internal protocols								
PF_INET	Internet Protocol Version 4 (IPv4)								
PF_INET6	Internet Protocol Version 6 (IPv6)								
PF_NCA	Network Cache and Accelerator (NCA) protocols								



must be specified in this manner. The protocol number to use is particular to the “communication domain” in which communication is to take place. If a protocol is specified by the caller, then it will be packaged into a socket level option request and sent to the underlying protocol layers.

Sockets of type `SOCK_STREAM` are full-duplex byte streams, similar to pipes. A stream socket must be in a *connected* state before any data may be sent or received on it. A connection to another socket is created with a `connect(3SOCKET)` call. Once connected, data may be transferred using `read(2)` and `write(2)` calls or some variant of the `send(3SOCKET)` and `recv(3SOCKET)` calls. When a session has been completed, a `close(2)` may be performed. Out-of-band data may also be transmitted as described on the `send(3SOCKET)` manual page and received as described on the `recv(3SOCKET)` manual page.

The communications protocols used to implement a `SOCK_STREAM` insure that data is not lost or duplicated. If a piece of data for which the peer protocol has buffer space cannot be successfully transmitted within a reasonable length of time, then the connection is considered broken and calls will indicate an error with `-1` returns and with `ETIMEDOUT` as the specific code in the global variable `errno`. The protocols optionally keep sockets “warm” by forcing transmissions roughly every minute in the absence of other activity. An error is then indicated if no response can be elicited on an otherwise idle connection for an extended period (for instance 5 minutes). A `SIGPIPE` signal is raised if a process sends on a broken stream; this causes naive processes, which do not handle the signal, to exit.

`SOCK_SEQPACKET` sockets employ the same system calls as `SOCK_STREAM` sockets. The only difference is that `read(2)` calls will return only the amount of data requested, and any remaining in the arriving packet will be discarded.

`SOCK_DGRAM` and `SOCK_RAW` sockets allow datagrams to be sent to correspondents named in `sendto(3SOCKET)` calls. Datagrams are generally received with `recvfrom(3SOCKET)`, which returns the next datagram with its return address.

An `fcntl(2)` call can be used to specify a process group to receive a `SIGURG` signal when the out-of-band data arrives. It may also enable non-blocking I/O and asynchronous notification of I/O events with `SIGIO` signals.

The operation of sockets is controlled by socket level *options*. These options are defined in the file `<sys/socket.h>`. `setsockopt(3SOCKET)` and `getsockopt(3SOCKET)` are used to set and get options, respectively.

## RETURN VALUES

A `-1` is returned if an error occurs. Otherwise the return value is a descriptor referencing the socket.

## ERRORS

The `socket()` call fails if:

<code>EACCES</code>	Permission to create a socket of the specified type and/or protocol is denied.
<code>EMFILE</code>	The per-process descriptor table is full.

socket(3SOCKET)

- ENOMEM Insufficient user memory is available.
- ENOSR There were insufficient STREAMS resources available to complete the operation.
- EPROTONOSUPPORT The protocol type or the specified protocol is not supported within this domain.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `nca(1)`, `close(2)`, `fcntl(2)`, `ioctl(2)`, `read(2)`, `write(2)`, `accept(3SOCKET)`, `bind(3SOCKET)`, `connect(3SOCKET)`, `getsockname(3SOCKET)`, `getsockopt(3SOCKET)`, `listen(3SOCKET)`, `recv(3SOCKET)`, `setsockopt(3SOCKET)`, `send(3SOCKET)`, `shutdown(3SOCKET)`, `socketpair(3SOCKET)`, `attributes(5)`, `in(3HEAD)`, `socket(3HEAD)`

<b>NAME</b>	socket – create an endpoint for communication																		
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -l<i>ynet</i> [ <i>library</i> ... ] #include &lt;sys/socket.h&gt;  int <b>socket</b>(int <i>domain</i>, int <i>type</i>, int <i>protocol</i>);</pre>																		
<b>DESCRIPTION</b>	<p>The <code>socket()</code> function creates an unbound socket in a communications domain, and returns a file descriptor that can be used in later function calls that operate on sockets.</p> <p>The function takes the following arguments:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><i>domain</i></td> <td>Specifies the communications domain in which a socket is to be created.</td> </tr> <tr> <td><i>type</i></td> <td>Specifies the type of socket to be created.</td> </tr> <tr> <td><i>protocol</i></td> <td>Specifies a particular protocol to be used with the socket. Specifying a <i>protocol</i> of 0 causes <code>socket()</code> to use an unspecified default protocol appropriate for the requested socket type.</td> </tr> </table> <p>The <i>domain</i> argument specifies the address family used in the communications domain. The address families supported by the system are implementation-dependent.</p> <p>The <code>&lt;sys/socket.h&gt;</code> header defines at least the following values for the <i>domain</i> argument:</p> <table border="0"> <tr> <td style="padding-right: 20px;">AF_UNIX</td> <td>File system pathnames.</td> </tr> <tr> <td>AF_INET</td> <td>Internet Protocol version 4 (IPv4) address.</td> </tr> <tr> <td>AF_INET6</td> <td>Internet Protocol version 6 (IPv6) address.</td> </tr> </table> <p>The <i>type</i> argument specifies the socket type, which determines the semantics of communication over the socket. The socket types supported by the system are implementation-dependent. Possible socket types include:</p> <table border="0"> <tr> <td style="padding-right: 20px;">SOCK_STREAM</td> <td>Provides sequenced, reliable, bidirectional, connection-mode byte streams, and may provide a transmission mechanism for out-of-band data.</td> </tr> <tr> <td>SOCK_DGRAM</td> <td>Provides datagrams, which are connectionless-mode, unreliable messages of fixed maximum length.</td> </tr> <tr> <td>SOCK_SEQPACKET</td> <td>Provides sequenced, reliable, bidirectional, connection-mode transmission path for records. A record can be sent using one or more output operations and received using one or more input operations, but a single operation never transfers part of more than one record. Record boundaries are visible to the receiver via the MSG_EOR flag.</td> </tr> </table>	<i>domain</i>	Specifies the communications domain in which a socket is to be created.	<i>type</i>	Specifies the type of socket to be created.	<i>protocol</i>	Specifies a particular protocol to be used with the socket. Specifying a <i>protocol</i> of 0 causes <code>socket()</code> to use an unspecified default protocol appropriate for the requested socket type.	AF_UNIX	File system pathnames.	AF_INET	Internet Protocol version 4 (IPv4) address.	AF_INET6	Internet Protocol version 6 (IPv6) address.	SOCK_STREAM	Provides sequenced, reliable, bidirectional, connection-mode byte streams, and may provide a transmission mechanism for out-of-band data.	SOCK_DGRAM	Provides datagrams, which are connectionless-mode, unreliable messages of fixed maximum length.	SOCK_SEQPACKET	Provides sequenced, reliable, bidirectional, connection-mode transmission path for records. A record can be sent using one or more output operations and received using one or more input operations, but a single operation never transfers part of more than one record. Record boundaries are visible to the receiver via the MSG_EOR flag.
<i>domain</i>	Specifies the communications domain in which a socket is to be created.																		
<i>type</i>	Specifies the type of socket to be created.																		
<i>protocol</i>	Specifies a particular protocol to be used with the socket. Specifying a <i>protocol</i> of 0 causes <code>socket()</code> to use an unspecified default protocol appropriate for the requested socket type.																		
AF_UNIX	File system pathnames.																		
AF_INET	Internet Protocol version 4 (IPv4) address.																		
AF_INET6	Internet Protocol version 6 (IPv6) address.																		
SOCK_STREAM	Provides sequenced, reliable, bidirectional, connection-mode byte streams, and may provide a transmission mechanism for out-of-band data.																		
SOCK_DGRAM	Provides datagrams, which are connectionless-mode, unreliable messages of fixed maximum length.																		
SOCK_SEQPACKET	Provides sequenced, reliable, bidirectional, connection-mode transmission path for records. A record can be sent using one or more output operations and received using one or more input operations, but a single operation never transfers part of more than one record. Record boundaries are visible to the receiver via the MSG_EOR flag.																		

## socket(3XNET)

If the *protocol* argument is non-zero, it must specify a protocol that is supported by the address family. The protocols supported by the system are implementation-dependent.

The process may need to have appropriate privileges to use the `socket ()` function or to create some sockets.

**USAGE** The documentation for specific address families specify which protocols each address family supports. The documentation for specific protocols specify which socket types each protocol supports.

The application can determine if an address family is supported by trying to create a socket with *domain* set to the protocol in question.

**RETURN VALUES** Upon successful completion, `socket ()` returns a nonnegative integer, the socket file descriptor. Otherwise a value of `-1` is returned and `errno` is set to indicate the error.

**ERRORS** The `socket ()` function will fail if:

EAFNOSUPPORT	The implementation does not support the specified address family.
EMFILE	No more file descriptors are available for this process.
ENFILE	No more file descriptors are available for the system.
EPROTONOSUPPORT	The protocol is not supported by the address family, or the protocol is not supported by the implementation.
EPROTOTYPE	The socket type is not supported by the protocol.

The `socket ()` function may fail if:

EACCES	The process does not have appropriate privileges.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
ENOMEM	Insufficient memory was available to fulfill the request.
ENOSR	There were insufficient STREAMS resources available for the operation to complete.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `accept(3XNET)`, `bind(3XNET)`, `connect(3XNET)`, `getsockname(3XNET)`, `getsockopt(3XNET)`, `listen(3XNET)`, `recv(3XNET)`, `recvfrom(3XNET)`, `recvmsg(3XNET)`, `send(3XNET)`, `sendmsg(3XNET)`, `setsockopt(3XNET)`, `shutdown(3XNET)`, `socketpair(3XNET)`, `attributes(5)`

<b>NAME</b>	socketpair – create a pair of connected sockets														
<b>SYNOPSIS</b>	<pre>cc [ <i>flag</i> ... ] <i>file</i> ... -lsocket -lnsl [ <i>library</i> ... ] #include &lt;sys/types.h&gt; #include &lt;sys/socket.h&gt;  int <b>socketpair</b>(int <i>domain</i>, int <i>type</i>, int <i>protocol</i>, int <i>sv</i>[2]);</pre>														
<b>DESCRIPTION</b>	The <code>socketpair()</code> library call creates an unnamed pair of connected sockets in the specified address family <i>domain</i> , of the specified <i>type</i> , that uses the optionally specified <i>protocol</i> . The descriptors that are used in referencing the new sockets are returned in <i>sv</i> [0] and <i>sv</i> [1]. The two sockets are indistinguishable.														
<b>RETURN VALUES</b>	<code>socketpair()</code> returns <code>-1</code> on failure and <code>0</code> on success.														
<b>ERRORS</b>	<p>The call succeeds unless:</p> <table border="0"> <tr> <td style="vertical-align: top;">EAFNOSUPPORT</td> <td>The specified address family is not supported on this machine.</td> </tr> <tr> <td style="vertical-align: top;">EMFILE</td> <td>Too many descriptors are in use by this process.</td> </tr> <tr> <td style="vertical-align: top;">ENOMEM</td> <td>There was insufficient user memory for the operation to complete.</td> </tr> <tr> <td style="vertical-align: top;">ENOSR</td> <td>There were insufficient STREAMS resources for the operation to complete.</td> </tr> <tr> <td style="vertical-align: top;">EOPNOTSUPP</td> <td>The specified protocol does not support creation of socket pairs.</td> </tr> <tr> <td style="vertical-align: top;">EPROTONOSUPPORT</td> <td>The specified protocol is not supported on this machine.</td> </tr> <tr> <td style="vertical-align: top;">EACCES</td> <td>The process does not have appropriate privileges.</td> </tr> </table>	EAFNOSUPPORT	The specified address family is not supported on this machine.	EMFILE	Too many descriptors are in use by this process.	ENOMEM	There was insufficient user memory for the operation to complete.	ENOSR	There were insufficient STREAMS resources for the operation to complete.	EOPNOTSUPP	The specified protocol does not support creation of socket pairs.	EPROTONOSUPPORT	The specified protocol is not supported on this machine.	EACCES	The process does not have appropriate privileges.
EAFNOSUPPORT	The specified address family is not supported on this machine.														
EMFILE	Too many descriptors are in use by this process.														
ENOMEM	There was insufficient user memory for the operation to complete.														
ENOSR	There were insufficient STREAMS resources for the operation to complete.														
EOPNOTSUPP	The specified protocol does not support creation of socket pairs.														
EPROTONOSUPPORT	The specified protocol is not supported on this machine.														
EACCES	The process does not have appropriate privileges.														
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:														
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe										
ATTRIBUTE TYPE	ATTRIBUTE VALUE														
MT-Level	Safe														
<b>SEE ALSO</b>	<code>pipe(2)</code> , <code>read(2)</code> , <code>write(2)</code> , <code>socket(3HEAD)</code> , <code>attributes(5)</code>														
<b>NOTES</b>	This call is currently implemented only for the <code>AF_UNIX</code> address family.														

## socketpair(3XNET)

<b>NAME</b>	socketpair – create a pair of connected sockets														
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lxnet [ library ... ] #include &lt;sys/socket.h&gt;  int socketpair(int domain, int type, int protocol, int socket_vector[2]);</pre>														
<b>DESCRIPTION</b>	<p>The <code>socketpair()</code> function creates an unbound pair of connected sockets in a specified <i>domain</i>, of a specified <i>type</i>, under the protocol optionally specified by the <i>protocol</i> argument. The two sockets are identical. The file descriptors used in referencing the created sockets are returned in <i>socket_vector0</i> and <i>socket_vector1</i>.</p> <table><tr><td><i>domain</i></td><td>Specifies the communications domain in which the sockets are to be created.</td></tr><tr><td><i>type</i></td><td>Specifies the type of sockets to be created.</td></tr><tr><td><i>protocol</i></td><td>Specifies a particular protocol to be used with the sockets. Specifying a <i>protocol</i> of 0 causes <code>socketpair()</code> to use an unspecified default protocol appropriate for the requested socket type.</td></tr><tr><td><i>socket_vector</i></td><td>Specifies a 2-integer array to hold the file descriptors of the created socket pair.</td></tr></table> <p>The <i>type</i> argument specifies the socket type, which determines the semantics of communications over the socket. The socket types supported by the system are implementation-dependent. Possible socket types include:</p> <table><tr><td>SOCK_STREAM</td><td>Provides sequenced, reliable, bidirectional, connection-mode byte streams, and may provide a transmission mechanism for out-of-band data.</td></tr><tr><td>SOCK_DGRAM</td><td>Provides datagrams, which are connectionless-mode, unreliable messages of fixed maximum length.</td></tr><tr><td>SOCK_SEQPACKET</td><td>Provides sequenced, reliable, bidirectional, connection-mode transmission path for records. A record can be sent using one or more output operations and received using one or more input operations, but a single operation never transfers part of more than one record. Record boundaries are visible to the receiver via the MSG_EOR flag.</td></tr></table> <p>If the <i>protocol</i> argument is non-zero, it must specify a protocol that is supported by the address family. The protocols supported by the system are implementation-dependent.</p> <p>The process may need to have appropriate privileges to use the <code>socketpair()</code> function or to create some sockets.</p>	<i>domain</i>	Specifies the communications domain in which the sockets are to be created.	<i>type</i>	Specifies the type of sockets to be created.	<i>protocol</i>	Specifies a particular protocol to be used with the sockets. Specifying a <i>protocol</i> of 0 causes <code>socketpair()</code> to use an unspecified default protocol appropriate for the requested socket type.	<i>socket_vector</i>	Specifies a 2-integer array to hold the file descriptors of the created socket pair.	SOCK_STREAM	Provides sequenced, reliable, bidirectional, connection-mode byte streams, and may provide a transmission mechanism for out-of-band data.	SOCK_DGRAM	Provides datagrams, which are connectionless-mode, unreliable messages of fixed maximum length.	SOCK_SEQPACKET	Provides sequenced, reliable, bidirectional, connection-mode transmission path for records. A record can be sent using one or more output operations and received using one or more input operations, but a single operation never transfers part of more than one record. Record boundaries are visible to the receiver via the MSG_EOR flag.
<i>domain</i>	Specifies the communications domain in which the sockets are to be created.														
<i>type</i>	Specifies the type of sockets to be created.														
<i>protocol</i>	Specifies a particular protocol to be used with the sockets. Specifying a <i>protocol</i> of 0 causes <code>socketpair()</code> to use an unspecified default protocol appropriate for the requested socket type.														
<i>socket_vector</i>	Specifies a 2-integer array to hold the file descriptors of the created socket pair.														
SOCK_STREAM	Provides sequenced, reliable, bidirectional, connection-mode byte streams, and may provide a transmission mechanism for out-of-band data.														
SOCK_DGRAM	Provides datagrams, which are connectionless-mode, unreliable messages of fixed maximum length.														
SOCK_SEQPACKET	Provides sequenced, reliable, bidirectional, connection-mode transmission path for records. A record can be sent using one or more output operations and received using one or more input operations, but a single operation never transfers part of more than one record. Record boundaries are visible to the receiver via the MSG_EOR flag.														

**USAGE** The documentation for specific address families specifies which protocols each address family supports. The documentation for specific protocols specifies which socket types each protocol supports.

The `socketpair()` function is used primarily with UNIX domain sockets and need not be supported for other domains.

**RETURN VALUES** Upon successful completion, this function returns 0. Otherwise, -1 is returned and `errno` is set to indicate the error.

**ERRORS** The `socketpair()` function will fail if:

- `EAFNOSUPPORT` The implementation does not support the specified address family.
- `EMFILE` No more file descriptors are available for this process.
- `ENFILE` No more file descriptors are available for the system.
- `EOPNOTSUPP` The specified protocol does not permit creation of socket pairs.
- `EPROTONOSUPPORT` The protocol is not supported by the address family, or the protocol is not supported by the implementation.
- `EPROTOTYPE` The socket type is not supported by the protocol.

The `socketpair()` function may fail if:

- `EACCES` The process does not have appropriate privileges.
- `ENOBUFS` Insufficient resources were available in the system to perform the operation.
- `ENOMEM` Insufficient memory was available to fulfill the request.
- `ENOSR` There were insufficient STREAMS resources available for the operation to complete.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `socket(3XNET)`, `attributes(5)`

## spray(3SOCKET)

<b>NAME</b>	spray – scatter data in order to test the network						
<b>SYNOPSIS</b>	<pre>cc [ flag ... ] file ... -lsocket -lnsl [ library ... ] #include &lt;rpcsvc/spray.h&gt;  bool_t xdr_sprayarr(XDR *xdrs, sprayarr *objp); bool_t xdr_spraycumul(XDR *xdrs, spraycumul *objp);</pre>						
<b>DESCRIPTION</b>	<p>The spray program sends packets to a given machine to test communications with that machine.</p> <p>The spray program is not a C function interface, per se, but it can be accessed using the generic remote procedure calling interface <code>clnt_call()</code>. See <code>rpc_clnt_calls(3NSL)</code>. The program sends a packet to the called host. The host acknowledges receipt of the packet. The program counts the number of acknowledgments and can return that count.</p> <p>The spray program currently supports the following procedures, which should be called in the order given:</p> <table><tr><td>SPRAYPROC_CLEAR</td><td>This procedure clears the counter.</td></tr><tr><td>SPRAYPROC_SPRAY</td><td>This procedure sends the packet.</td></tr><tr><td>SPRAYPROC_GET</td><td>This procedure returns the count and the amount of time since the last SPRAYPROC_CLEAR.</td></tr></table>	SPRAYPROC_CLEAR	This procedure clears the counter.	SPRAYPROC_SPRAY	This procedure sends the packet.	SPRAYPROC_GET	This procedure returns the count and the amount of time since the last SPRAYPROC_CLEAR.
SPRAYPROC_CLEAR	This procedure clears the counter.						
SPRAYPROC_SPRAY	This procedure sends the packet.						
SPRAYPROC_GET	This procedure returns the count and the amount of time since the last SPRAYPROC_CLEAR.						
<b>EXAMPLES</b>	<p><b>EXAMPLE 1</b> Using <code>spray()</code></p> <p>The following code fragment demonstrates how the spray program is used:</p> <pre>#include &lt;rpc/rpc.h&gt; #include &lt;rpcsvc/spray.h&gt; . . . spraycumul    spray_result; sprayarr      spray_data; char          buf[100];          /* arbitrary data */ int           loop = 1000; CLIENT       *clnt; struct timeval timeout0 = {0, 0}; struct timeval timeout25 = {25, 0}; spray_data.sprayarr_len = (uint_t)100; spray_data.sprayarr_val = buf; clnt = clnt_create("somehost", SPRAYPROC, SPRAYVERS, "netpath"); if (clnt == (CLIENT *)NULL) {     /* handle this error */ } if (clnt_call(clnt, SPRAYPROC_CLEAR,              xdr_void, NULL, xdr_void, NULL, timeout25)) {     /* handle this error */ } while (loop- &gt; 0) {     if (clnt_call(clnt, SPRAYPROC_SPRAY,                  xdr_sprayarr, &amp;spray_data, xdr_void, NULL, timeout0)) {         /* handle this error */     } }</pre>						



**EXAMPLE 1** Using `spray()` (Continued)

```

    }
}
if (clnt_call(clnt, SPRAYPROC_GET,
             xdr_void, NULL, xdr_spraycumul, &spray_result, timeout25)) {
    /* handle this error */
}
printf("Acknowledged %ld of 1000 packets in %d secs %d usecs\n",
       spray_result.counter,
       spray_result.clock.sec,
       spray_result.clock.usec);

```

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

**SEE ALSO** `spray(1M)`, `rpc_clnt_calls(3NSL)`, `attributes(5)`

**NOTES** This interface is unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.

A `spray` program is not useful as a networking benchmark as it uses unreliable connectionless transports, for example, `udp`. It can report a large number of packets dropped, when the drops were caused by the program sending packets faster than they can be buffered locally, that is, before the packets get to the network medium.

## t\_accept(3NSL)

<b>NAME</b>	t_accept – accept a connection request
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_accept(int fd, int resfd, const struct t_call *call);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces that evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, a different header file, <code>tuser.h</code>, must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function is issued by a transport user to accept a connection request. The parameter <i>fd</i> identifies the local transport endpoint where the connection indication arrived; <i>resfd</i> specifies the local transport endpoint where the connection is to be established, and <i>call</i> contains information required by the transport provider to complete the connection. The parameter <i>call</i> points to a <code>t_call</code> structure which contains the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata; int sequence;</pre> <p>In <i>call</i>, <i>addr</i> is the protocol address of the calling transport user, <i>opt</i> indicates any options associated with the connection, <i>udata</i> points to any user data to be returned to the caller, and <i>sequence</i> is the value returned by <code>t_listen(3NSL)</code> that uniquely associates the response with a previously received connection indication. The address of the caller, <i>addr</i> may be null (length zero). Where <i>addr</i> is not null then it may optionally be checked by XTI.</p> <p>A transport user may accept a connection on either the same, or on a different, local transport endpoint than the one on which the connection indication arrived. Before the connection can be accepted on the same endpoint (<i>resfd==fd</i>), the user must have responded to any previous connection indications received on that transport endpoint by means of <code>t_accept()</code> or <code>t_snddis(3NSL)</code>. Otherwise, <code>t_accept()</code> will fail and set <code>t_errno</code> to <code>TINDOUT</code>.</p> <p>If a different transport endpoint is specified (<i>resfd!=fd</i>), then the user may or may not choose to bind the endpoint before the <code>t_accept()</code> is issued. If the endpoint is not bound prior to the <code>t_accept()</code>, the endpoint must be in the <code>T_UNBND</code> state before the <code>t_accept()</code> is issued, and the transport provider will automatically bind it to an address that is appropriate for the protocol concerned. If the transport user chooses to bind the endpoint it must be bound to a protocol address with a <i>qlen</i> of zero and must be in the <code>T_IDLE</code> state before the <code>t_accept()</code> is issued.</p> <p>Responding endpoints should be supplied to <code>t_accept()</code> in the state <code>T_UNBND</code>.</p> <p>The call to <code>t_accept()</code> may fail with <code>t_errno</code> set to <code>TLOOK</code> if there are indications (for example connect or disconnect) waiting to be received on endpoint <i>fd</i>. Applications should be prepared for such a failure.</p>

t\_accept(3NSL)

The *udata* argument enables the called transport user to send user data to the caller and the amount of user data must not exceed the limits supported by the transport provider as returned in the *connect* field of the *info* argument of t\_open(3NSL) or t\_getinfo(3NSL). If the *len* field of *udata* is zero, no data will be sent to the caller. All the *maxlen* fields are meaningless.

When the user does not indicate any option (*call*→*opt.len* = 0) the connection shall be accepted with the option values currently set for the responding endpoint *resfd*.

**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.

**VALID STATES** fd: T\_INCON  
resfd (fd!=resfd): T\_IDLE, T\_UNBND

**ERRORS** On failure, t\_errno is set to one of the following:

TACCES	The user does not have permission to accept a connection on the responding transport endpoint or to use the specified options.
TBADADDR	The specified protocol address was in an incorrect format or contained illegal information.
TBADDATA	The amount of user data specified was not within the bounds allowed by the transport provider.
TBADF	The file descriptor <i>fd</i> or <i>resfd</i> does not refer to a transport endpoint.
TBADOPT	The specified options were in an incorrect format or contained illegal information.
TBADSEQ	Either an invalid sequence number was specified, or a valid sequence number was specified but the connection request was aborted by the peer. In the latter case, its T_DISCONNECT event will be received on the listening endpoint.
TINDOUT	The function was called with <i>fd</i> == <i>resfd</i> but there are outstanding connection indications on the endpoint. Those other connection indications must be handled either by rejecting them by means of t_snddis(3NSL) or accepting them on a different endpoint by means of t_accept.
TLOOK	An asynchronous event has occurred on the transport endpoint referenced by <i>fd</i> and requires immediate attention.
TNOTSUPPORT	This function is not supported by the underlying transport provider.

## t\_accept(3NSL)

	TOUTSTATE	The communications endpoint referenced by <i>fd</i> or <i>resfd</i> is not in one of the states in which a call to this function is valid.
	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <i>t_errno</i> ).
	TPROVMISMATCH	The file descriptors <i>fd</i> and <i>resfd</i> do not refer to the same transport provider.
	TRESADDR	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.
	TRESQLEN	The endpoint referenced by <i>resfd</i> (where <i>resfd</i> != <i>fd</i> ) was bound to a protocol address with a <i>qlen</i> that is greater than zero.
	TSYSERR	A system error has occurred during execution of this function.
<b>TLI COMPATIBILITY</b>		The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>		The XTI interfaces use the header file, <i>xti.h</i> . TLI interfaces should <i>not</i> use this header. They should use the header: <pre>#include &lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>		The <i>t_errno</i> values that can be set by the XTI interface and cannot be set by the TLI interface are:  TPROTO TINDOUT TPROVMISMATCH TRESADDR TRESQLEN
<b>Option Buffer</b>		The format of the options in an <i>opt</i> buffer is dictated by the transport provider. Unlike the XTI interface, the TLI interface does not specify the buffer format.
<b>ATTRIBUTES</b>		See <i>attributes(5)</i> for descriptions of the following attributes:

t\_accept(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** t\_connect(3NSL), t\_getinfo(3NSL), t\_getstate(3NSL), t\_listen(3NSL), t\_open(3NSL), t\_optmgmt(3NSL), t\_rcvconnect(3NSL), t\_snddis(3NSL), attributes(5)

**WARNINGS** There may be transport provider-specific restrictions on address binding.

Some transport providers do not differentiate between a connection indication and the connection itself. If the connection has already been established after a successful return of t\_listen(3NSL), t\_accept() will assign the existing connection to the transport endpoint specified by *resfd*.

## t\_alloc(3NSL)

<b>NAME</b>	t_alloc – allocate a library structure																													
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  void *t_alloc(int fd, int struct_type, int fields);</pre>																													
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, a different header file, <code>tiuser.h</code>, must be used. Refer to the section, TLI COMPATIBILITY, for a description of differences between the two interfaces.</p> <p>The <code>t_alloc()</code> function dynamically allocates memory for the various transport function argument structures as specified below. This function will allocate memory for the specified structure, and will also allocate memory for buffers referenced by the structure.</p> <p>The structure to allocate is specified by <code>struct_type</code> and must be one of the following:</p> <table><tr><td>T_BIND</td><td>struct</td><td>t_bind</td></tr><tr><td>T_CALL</td><td>struct</td><td>t_call</td></tr><tr><td>T_OPTMGMT</td><td>struct</td><td>t_optmgmt</td></tr><tr><td>T_DIS</td><td>struct</td><td>t_discon</td></tr><tr><td>T_UNITDATA</td><td>struct</td><td>t_unitdata</td></tr><tr><td>T_UDERROR</td><td>struct</td><td>t_uderr</td></tr><tr><td>T_INFO</td><td>struct</td><td>t_info</td></tr></table> <p>where each of these structures may subsequently be used as an argument to one or more transport functions.</p> <p>Each of the above structures, except T_INFO, contains at least one field of type <code>struct netbuf</code>. 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 <code>info</code> argument of <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>. The relevant fields of the <code>info</code> argument are described in the following list. The <code>fields</code> argument specifies which buffers to allocate, where the argument is the bitwise-or of any of the following:</p> <table><tr><td>T_ADDR</td><td>The <code>addr</code> field of the <code>t_bind</code>, <code>t_call</code>, <code>t_unitdata</code> or <code>t_uderr</code> structures.</td></tr><tr><td>T_OPT</td><td>The <code>opt</code> field of the <code>t_optmgmt</code>, <code>t_call</code>, <code>t_unitdata</code> or <code>t_uderr</code> structures.</td></tr><tr><td>T_UDATA</td><td>The <code>udata</code> field of the <code>t_call</code>, <code>t_discon</code> or <code>t_unitdata</code> structures.</td></tr><tr><td>T_ALL</td><td>All relevant fields of the given structure. Fields which are not supported by the transport provider specified by <code>fd</code> will not be allocated.</td></tr></table>	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	T_ADDR	The <code>addr</code> field of the <code>t_bind</code> , <code>t_call</code> , <code>t_unitdata</code> or <code>t_uderr</code> structures.	T_OPT	The <code>opt</code> field of the <code>t_optmgmt</code> , <code>t_call</code> , <code>t_unitdata</code> or <code>t_uderr</code> structures.	T_UDATA	The <code>udata</code> field of the <code>t_call</code> , <code>t_discon</code> or <code>t_unitdata</code> structures.	T_ALL	All relevant fields of the given structure. Fields which are not supported by the transport provider specified by <code>fd</code> will not be allocated.
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																												
T_ADDR	The <code>addr</code> field of the <code>t_bind</code> , <code>t_call</code> , <code>t_unitdata</code> or <code>t_uderr</code> structures.																													
T_OPT	The <code>opt</code> field of the <code>t_optmgmt</code> , <code>t_call</code> , <code>t_unitdata</code> or <code>t_uderr</code> structures.																													
T_UDATA	The <code>udata</code> field of the <code>t_call</code> , <code>t_discon</code> or <code>t_unitdata</code> structures.																													
T_ALL	All relevant fields of the given structure. Fields which are not supported by the transport provider specified by <code>fd</code> will not be allocated.																													

For each relevant field specified in *fields*, `t_alloc()` will allocate memory for the buffer associated with the field, and initialize the *len* field to zero and the *buf* pointer and *maxlen* field accordingly. Irrelevant or unknown values passed in fields are ignored. Since the length of the buffer allocated will be based on the same size information that is returned to the user on a call to `t_open(3NSL)` and `t_getinfo(3NSL)`, *fd* must refer to the transport endpoint through which the newly allocated structure will be passed. In the case where a `T_INFO` structure is to be allocated, *fd* may be set to any value. In this way the appropriate size information can be accessed. If the size value associated with any specified field is `T_INVALID`, `t_alloc()` will be unable to determine the size of the buffer to allocate and will fail, setting `t_errno` to `TSYSERR` and `errno` to `EINVAL`. See `t_open(3NSL)` or `t_getinfo(3NSL)`. If the size value associated with any specified field is `T_INFINITE`, then the behavior of `t_alloc()` is implementation-defined. For any field not specified in *fields*, *buf* will be set to the null pointer and *len* and *maxlen* will be set to zero. See `t_open(3NSL)` or `t_getinfo(3NSL)`.

The pointer returned if the allocation succeeds is suitably aligned so that it can be assigned to a pointer to any type of object and then used to access such an object or array of such objects in the space allocated.

Use of `t_alloc()` to allocate structures will help ensure the compatibility of user programs with future releases of the transport interface functions.

**RETURN VALUES**

On successful completion, `t_alloc()` returns a pointer to the newly allocated structure. On failure, a null pointer is returned.

**VALID STATES**

ALL - apart from `T_UNINIT`

**ERRORS**

On failure, `t_errno` is set to one of the following:

<code>TBADF</code>	<code>struct_type</code> is other than <code>T_INFO</code> and the specified file descriptor does not refer to a transport endpoint.
<code>TNOSTRUCTYPE</code>	Unsupported <code>struct_type</code> requested. This can include a request for a structure type which is inconsistent with the transport provider type specified, that is, connection-mode or connectionless-mode.
<code>TPROTO</code>	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
<code>TSYSERR</code>	A system error has occurred during execution of this function.

**TLI COMPATIBILITY**

The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header**

The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

## t\_alloc(3NSL)

### Error Description Values

The `t_errno` values that can be set by the XTI interface and cannot be set by the TLI interface are:

`TPROTO`

`TNOSTRUCTYPE`

### Special Buffer Sizes

Assume that the value associated with any field of `struct t_info` (argument returned by `t_open()` or `t_getinfo()`) that describes buffer limits is `-1`. Then the underlying service provider can support a buffer of unlimited size. If this is the case, `t_alloc()` will allocate a buffer with the default size 1024 bytes, which may be handled as described in the next paragraph.

If the underlying service provider supports a buffer of unlimited size in the `netbuf` structure (see `t_connect(3NSL)`), `t_alloc()` will return a buffer of size 1024 bytes. If a larger size buffer is required, it will need to be allocated separately using a memory allocation routine such as `malloc(3C)`. The `buf` and `maxlen` fields of the `netbuf` data structure can then be updated with the address of the new buffer and the 1024 byte buffer originally allocated by `t_alloc()` can be freed using `free(3C)`.

Assume that the value associated with any field of `struct t_info` (argument returned by `t_open()` or `t_getinfo()`) that describes `nbuffer` limits is `-2`. Then `t_alloc()` will set the buffer pointer to `NULL` and the buffer maximum size to 0, and then will return success (see `t_open(3NSL)` or `t_getinfo(3NSL)`).

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

### SEE ALSO

`free(3C)`, `malloc(3C)`, `t_connect(3NSL)`, `t_free(3NSL)`, `t_getinfo(3NSL)`, `t_open(3NSL)`, `attributes(5)`



<b>NAME</b>	t_bind – bind an address to a transport endpoint
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_bind(int fd, const struct t_bind *req, struct t_bind *ret);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces that evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function associates a protocol address with the transport endpoint specified by <i>fd</i> 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.</p> <p>The <i>req</i> and <i>ret</i> arguments point to a <code>t_bind</code> structure containing the following members:</p> <pre>struct netbuf    addr; unsigned        qlen;</pre> <p>The <i>addr</i> field of the <code>t_bind</code> structure specifies a protocol address, and the <i>qlen</i> field is used to indicate the maximum number of outstanding connection indications.</p> <p>The parameter <i>req</i> is used to request that an address, represented by the <code>netbuf</code> structure, be bound to the given transport endpoint. The parameter <i>len</i> specifies the number of bytes in the address, and <i>buf</i> points to the address buffer. The parameter <i>maxlen</i> has no meaning for the <i>req</i> argument. On return, <i>ret</i> contains an encoding for the address that the transport provider actually bound to the transport endpoint; if an address was specified in <i>req</i>, this will be an encoding of the same address. In <i>ret</i>, the user specifies <i>maxlen</i>, which is the maximum size of the address buffer, and <i>buf</i> which points to the buffer where the address is to be placed. On return, <i>len</i> specifies the number of bytes in the bound address, and <i>buf</i> points to the bound address. If <i>maxlen</i> equals zero, no address is returned. If <i>maxlen</i> is greater than zero and less than the length of the address, <code>t_bind()</code> fails with <code>t_errno</code> set to <code>TBUFOVFLW</code>.</p> <p>If the requested address is not available, <code>t_bind()</code> will return <code>-1</code> with <code>t_errno</code> set as appropriate. If no address is specified in <i>req</i> (the <i>len</i> field of <i>addr</i> in <i>req</i> is zero or <i>req</i> is <code>NULL</code>), the transport provider will assign an appropriate address to be bound, and will return that address in the <i>addr</i> field of <i>ret</i>. If the transport provider could not allocate an address, <code>t_bind()</code> will fail with <code>t_errno</code> set to <code>TNOADDR</code>.</p> <p>The parameter <i>req</i> may be a null pointer if the user does not wish to specify an address to be bound. Here, the value of <i>qlen</i> is assumed to be zero, and the transport provider will assign an address to the transport endpoint. Similarly, <i>ret</i> may be a null pointer if the user does not care what address was bound by the provider and is not interested</p>

## t\_bind(3NSL)

in the negotiated value of *qlen*. It is valid to set *req* and *ret* 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.

The *qlen* field has meaning only when initializing a connection-mode service. It specifies the number of outstanding connection indications that the transport provider should support for the given transport endpoint. An outstanding connection indication is one that has been passed to the transport user by the transport provider but which has not been accepted or rejected. A value of *qlen* greater than zero is only meaningful when issued by a passive transport user that expects other users to call it. The value of *qlen* will be negotiated by the transport provider and may be changed if the transport provider cannot support the specified number of outstanding connection indications. However, this value of *qlen* will never be negotiated from a requested value greater than zero to zero. This is a requirement on transport providers; see WARNINGS below. On return, the *qlen* field in *ret* will contain the negotiated value.

If *fd* refers to a connection-mode service, this function allows more than one transport endpoint to be bound to the same protocol address. but it is not possible to bind more than one protocol address to the same transport endpoint. However, the transport provider must also support this capability. If a user binds more than one transport endpoint to the same protocol address, only one endpoint can be used to listen for connection indications associated with that protocol address. In other words, only one `t_bind()` for a given protocol address may specify a value of *qlen* greater than zero. In this way, the transport provider can identify which transport endpoint should be notified of an incoming connection indication. If a user attempts to bind a protocol address to a second transport endpoint with a value of *qlen* 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(3NSL)` or `t_close(3NSL)` 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 connection indications.

If *fd* refers to connectionless mode service, this function allows for more than one transport endpoint to be associated with a protocol address, where the underlying transport provider supports this capability (often in conjunction with value of a protocol-specific option). If a user attempts to bind a second transport endpoint to an already bound protocol address when such capability is not supported for a transport provider, `t_bind()` will return `-1` and set `t_errno` to `TADDRBUSY`.

<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of <code>-1</code> is returned and <code>t_errno</code> is set to indicate an error.		
<b>VALID STATES</b>	<code>T_UNBND</code>		
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following: <table><tr><td><code>TACCES</code></td><td>The user does not have permission to use the specified address.</td></tr></table>	<code>TACCES</code>	The user does not have permission to use the specified address.
<code>TACCES</code>	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.
TBUFOVFLW	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 T_IDLE and the information to be returned in <i>ret</i> will be discarded.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TNOADDR	The transport provider could not allocate an address.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <i>t_errno</i> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI  
COMPATIBILITY**

The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header**

The XTI interfaces use the header file, *xti.h*. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Address Bound**

The user can compare the addresses in *req* and *ret* to determine whether the transport provider bound the transport endpoint to a different address than that requested.

**Error Description  
Values**

The *t\_errno* values TPROTO and TADDRBUSY can be set by the XTI interface but cannot be set by the TLI interface.

A *t\_errno* value that this routine can return under different circumstances than its XTI counterpart is TBUFOVFLW. It can be returned even when the *maxlen* field of the corresponding buffer has been set to zero.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**

*t\_accept(3NSL)*, *t\_alloc(3NSL)*, *t\_close(3NSL)*, *t\_connect(3NSL)*, *t\_unbind(3NSL)*, *attributes(5)*

t\_bind(3NSL)

**WARNINGS**

The requirement that the value of *qlen* never be negotiated from a requested value greater than zero to zero implies that transport providers, rather than the XTI implementation itself, accept this restriction.

An implementation need not allow an application explicitly to bind more than one communications endpoint to a single protocol address, while permitting more than one connection to be accepted to the same protocol address. That means that although an attempt to bind a communications endpoint to some address with *qlen=0* might be rejected with `TADDRBUSY`, the user may nevertheless use this (unbound) endpoint as a responding endpoint in a call to `t_accept(3NSL)`. To become independent of such implementation differences, the user should supply unbound responding endpoints to `t_accept(3NSL)`.

The local address bound to an endpoint may change as result of a `t_accept(3NSL)` or `t_connect(3NSL)` call. Such changes are not necessarily reversed when the connection is released.

<b>NAME</b>	t_close – close a transport endpoint
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_close(int fd);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>The <code>t_close()</code> function informs the transport provider that the user is finished with the transport endpoint specified by <code>fd</code>, and frees any local library resources associated with the endpoint. In addition, <code>t_close()</code> closes the file associated with the transport endpoint.</p> <p>The function <code>t_close()</code> should be called from the <code>T_UNBND</code> state. See <code>t_getstate(3NSL)</code>. However, this function 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, <code>close(2)</code> will be issued for that file descriptor; if there are no other descriptors in this process or in another process which references the communication endpoint, any connection that may be associated with that endpoint is broken. The connection may be terminated in an orderly or abortive manner.</p> <p>A <code>t_close()</code> 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.</p>
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.
<b>VALID STATES</b>	<code>T_UNBND</code>
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to the following: <ul style="list-style-type: none"> <li><code>TBADFD</code>           The specified file descriptor does not refer to a transport endpoint.</li> <li><code>TPROTO</code>           This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<code>t_errno</code>).</li> <li><code>TSYSERR</code>          A system error has occurred during execution of this function.</li> </ul>
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>	The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header:

t\_close(3NSL)

#include <tiuser.h>

**Error Description Values**

The t\_errno value that can be set by the XTI interface and cannot be set by the TLI interface is:

TPROTO

**ATTRIBUTES**

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**

close(2), t\_getstate(3NSL), t\_open(3NSL), t\_unbind(3NSL), attributes(5)

<b>NAME</b>	t_connect – establish a connection with another transport user
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_connect(int fd, const struct t_call *sndcall, struct t_call               *rcvcall);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces. This function enables a transport user to request a connection to the specified destination transport user.</p> <p>This function can only be issued in the <code>T_IDLE</code> state. The parameter <code>fd</code> identifies the local transport endpoint where communication will be established, while <code>sndcall</code> and <code>rcvcall</code> point to a <code>t_call</code> structure which contains the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata; int sequence;</pre> <p>The parameter <code>sndcall</code> specifies information needed by the transport provider to establish a connection and <code>rcvcall</code> specifies information that is associated with the newly established connection.</p> <p>In <code>sndcall</code>, <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>udata</code> points to optional user data that may be passed to the destination transport user during connection establishment, and <code>sequence</code> has no meaning for this function.</p> <p>On return, in <code>rcvcall</code>, <code>addr</code> contains the protocol address associated with the responding transport endpoint, <code>opt</code> represents any protocol-specific information associated with the connection, <code>udata</code> points to optional user data that may be returned by the destination transport user during connection establishment, and <code>sequence</code> has no meaning for this function.</p> <p>The <code>opt</code> argument permits users to define the options that may be passed to the transport provider. The user may choose not to negotiate protocol options by setting the <code>len</code> field of <code>opt</code> to zero. In this case, the provider uses the option values currently set for the communications endpoint.</p> <p>If used, <code>sndcall→opt.buf</code> must point to a buffer with the corresponding options, and <code>sndcall→opt.len</code> must specify its length. The <code>maxlen</code> and <code>buf</code> fields of the <code>netbuf</code> structure pointed to by <code>rcvcall→addr</code> and <code>rcvcall→opt</code> must be set before the call.</p>

## t\_connect(3NSL)

The *udata* argument 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* argument of `t_open(3NSL)` or `t_getinfo(3NSL)`. If the *len* of *udata* is zero in *sndcall*, no data will be sent to the destination transport user.

On return, the *addr*, *opt* and *udata* fields of *rcvcall* will be updated to reflect values associated with the connection. Thus, the *maxlen* field of each argument must be set before issuing this function to indicate the maximum size of the buffer for each. However, *maxlen* can be set to zero, in which case no information to this specific argument is given to the user on the return from `t_connect()`. If *maxlen* is greater than zero and less than the length of the value, `t_connect()` fails with `t_errno` set to `TBUFOVFLW`. If *rcvcall* is set to `NULL`, no information at all is returned.

By default, `t_connect()` executes in synchronous mode, and will wait for the destination user's response before returning control to the local user. A successful return (that is, return value of zero) indicates that the requested connection has been established. However, if `O_NONBLOCK` is set by means of `t_open(3NSL)` or `fcntl(2)`, `t_connect()` executes in asynchronous mode. In this case, the call will not wait for the remote user's response, but will return control immediately to the local user and return `-1` with `t_errno` 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 connection request to the destination transport user. The `t_rcvconnect(3NSL)` function is used in conjunction with `t_connect()` 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 `t_rcvconnect(3NSL)`, `t_rcvdis(3NSL)` or `t_snddis(3NSL)`. When an asynchronous `t_connect()` call is interrupted by the arrival of a signal, the state of the corresponding transport endpoint is `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.

**VALID STATES** `T_IDLE`

**ERRORS** On failure, `t_errno` is set to one of the following:

<code>TACCES</code>	The user does not have permission to use the specified address or options.
<code>TADDRBUSY</code>	This transport provider does not support multiple connections with the same local and remote addresses. This error indicates that a connection already exists.
<code>TBADADDR</code>	The specified protocol address was in an incorrect format or contained illegal information.



	TBADDATA	The amount of user data specified was not within the bounds allowed by the transport provider.
	TBADF	The specified file descriptor does not refer to a transport endpoint.
	TBADOPT	The specified protocol options were in an incorrect format or contained illegal information.
	TBUFOVFLW	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. If executed in synchronous mode, the provider's state, as seen by the user, changes to T_DATAXFER, and the information to be returned in <i>rcvcall</i> is discarded.
	TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
	TNODATA	O_NONBLOCK was set, so the function successfully initiated the connection establishment procedure, but did not wait for a response from the remote user.
	TNOTSUPPORT	This function is not supported by the underlying transport provider.
	TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <i>t_errno</i> ).
	TSYSERR	A system error has occurred during execution of this function.
<b>TLI COMPATIBILITY</b>		The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>		The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header:  <pre>#include &lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>		The TPROTO and TADDRBUSY <i>t_errno</i> values can be set by the XTI interface but not by the TLI interface.  A <i>t_errno</i> value that this routine can return under different circumstances than its XTI counterpart is TBUFOVFLW. It can be returned even when the <i>maxlen</i> field of the corresponding buffer has been set to zero.
<b>Option Buffers</b>		The format of the options in an <i>opt</i> buffer is dictated by the transport provider. Unlike the XTI interface, the TLI interface does not fix the buffer format.

t\_connect(3NSL)

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** fcntl(2), t\_accept(3NSL), t\_alloc(3NSL), t\_getinfo(3NSL), t\_listen(3NSL), t\_open(3NSL), t\_optmgmt(3NSL), t\_rcvconnect(3NSL), t\_rcvdis(3NSL), t\_snddis(3NSL), attributes

<b>NAME</b>	t_errno – XTI error return value
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;</pre>
<b>DESCRIPTION</b>	<p>This error return value is part of the XTI interfaces that evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI interface that has the same name as an XTI interfaces, a different headerfile, &lt;tiuser.h&gt;, must be used. Refer the the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>t_errno is used by XTI functions to return error values.</p> <p>XTI functions provide an error number in t_errno which has type <i>int</i> and is defined in &lt;xti.h&gt;. The value of t_errno will be defined only after a call to a XTI function for which it is explicitly stated to be set and until it is changed by the next XTI function call. The value of t_errno should only be examined when it is indicated to be valid by a function's return value. Programs should obtain the definition of t_errno by the inclusion of &lt;xti.h&gt;. The practice of defining t_errno in program as <code>extern int t_errno</code> is obsolescent. No XTI function sets t_errno to 0 to indicate an error.</p> <p>It is unspecified whether t_errno is a macro or an identifier with external linkage. It represents a modifiable lvalue of type <i>int</i>. If a macro definition is suppressed in order to access an actual object or a program defines an identifier with name <i>t_errno</i>, the behavior is undefined.</p> <p>The symbolic values stored in t_errno by an XTI function are defined in the ERRORS sections in all relevant XTI function definition pages.</p>
<b>TLI COMPATIBILITY</b>	<p>t_errno is also used by TLI functions to return error values.</p> <p>The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.</p>
<b>Interface Header</b>	<p>The XTI interfaces use the header file, &lt;xti.h&gt;. TLI interfaces should <i>not</i> use this header. They should use the header:</p> <pre>#include &lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>	<p>The t_errno values that can be set by the XTI interface but cannot be set by the TLI interface are:</p> <pre>TNOSTRUCTYPE TBADNAME TBADQLEN TADDRBUSY TINDOUT</pre>

t\_errno(3NSL)

TPROVMISMATCH  
TRESADDR  
TQFULL  
TPROTO

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** `attributes(5)`

<b>NAME</b>	t_error – produce error message
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_error(const char *errmsg);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>The <code>t_error()</code> function produces a message on the standard error output which describes the last error encountered during a call to a transport function. The argument string <code>errmsg</code> is a user-supplied error message that gives context to the error.</p> <p>The error message is written as follows: first (if <code>errmsg</code> is not a null pointer and the character pointed to be <code>errmsg</code> is not the null character) the string pointed to by <code>errmsg</code> followed by a colon and a space; then a standard error message string for the current error defined in <code>t_errno</code>. If <code>t_errno</code> has a value different from <code>TSYSERR</code>, the standard error message string is followed by a newline character. If, however, <code>t_errno</code> is equal to <code>TSYSERR</code>, the <code>t_errno</code> string is followed by the standard error message string for the current error defined in <code>errno</code> followed by a newline.</p> <p>The language for error message strings written by <code>t_error()</code> is that of the current locale. If it is English, the error message string describing the value in <code>t_errno</code> may be derived from the comments following the <code>t_errno</code> codes defined in <code>xti.h</code>. The contents of the error message strings describing the value in <code>errno</code> are the same as those returned by the <code>strerror(3C)</code> function with an argument of <code>errno</code>.</p> <p>The error number, <code>t_errno</code>, is only set when an error occurs and it is not cleared on successful calls.</p>
<b>EXAMPLES</b>	<p>If a <code>t_connect(3NSL)</code> function fails on transport endpoint <code>fd2</code> because a bad address was given, the following call might follow the failure:</p> <pre>t_error("t_connect failed on fd2");</pre> <p>The diagnostic message to be printed would look like:</p> <pre>t_connect failed on fd2: incorrect addr format</pre> <p>Where <i>incorrect addr format</i> identifies the specific error that occurred, and <i>t_connect failed on fd2</i> tells the user which function failed on which transport endpoint.</p>
<b>RETURN VALUES</b>	Upon completion, a value of 0 is returned.
<b>VALID STATES</b>	All - apart from <code>T_UNINIT</code>
<b>ERRORS</b>	No errors are defined for the <code>t_error()</code> function.

t\_error(3NSL)

**TLI  
COMPATIBILITY**

The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header**

The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Error Description  
Values**

The `t_errno` value that can be set by the XTI interface and cannot be set by the TLI interface is:

TPROTO

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**

`t_errno(3NSL)`, `strerror(3C)`, `attributes(5)`

<b>NAME</b>	t_free – free a library structure				
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_free(void *ptr, int struct_type);</pre>				
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>The <code>t_free()</code> function frees memory previously allocated by <code>t_alloc(3NSL)</code>. This function will free memory for the specified structure, and will also free memory for buffers referenced by the structure.</p> <p>The argument <code>ptr</code> points to one of the seven structure types described for <code>t_alloc(3NSL)</code>, and <code>struct_type</code> identifies the type of that structure which must be one of the following:</p> <pre>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</pre> <p>where each of these structures is used as an argument to one or more transport functions.</p> <p>The function <code>t_free()</code> will check the <code>addr</code>, <code>opt</code> and <code>udata</code> fields of the given structure, as appropriate, and free the buffers pointed to by the <code>buf</code> field of the <code>netbuf</code> structure. If <code>buf</code> is a null pointer, <code>t_free()</code> will not attempt to free memory. After all buffers are freed, <code>t_free()</code> will free the memory associated with the structure pointed to by <code>ptr</code>.</p> <p>Undefined results will occur if <code>ptr</code> or any of the <code>buf</code> pointers points to a block of memory that was not previously allocated by <code>t_alloc(3NSL)</code>.</p>				
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.				
<b>VALID STATES</b>	ALL - apart from <code>T_UNINIT</code> .				
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to the following:				
	<table border="0"> <tr> <td style="padding-right: 20px;">TNOSTRUCTYPE</td> <td>Unsupported <code>struct_type</code> requested.</td> </tr> <tr> <td>TPROTO</td> <td>This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<code>t_errno</code>).</td> </tr> </table>	TNOSTRUCTYPE	Unsupported <code>struct_type</code> requested.	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TNOSTRUCTYPE	Unsupported <code>struct_type</code> requested.				
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).				

t\_free(3NSL)

TSYSERR                      A system error has occurred during execution of this function.

**TLI COMPATIBILITY**      The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header**      The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Error Description Values**      The `t_errno` value that can be set by the XTI interface and cannot be set by the TLI interface is:

TPROTO

**ATTRIBUTES**              See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**                `t_alloc(3NSL)`, `attributes(5)`



<b>NAME</b>	t_getinfo – get protocol-specific service information
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_getinfo(int fd, struct t_info *info);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function returns the current characteristics of the underlying transport protocol and/or transport connection associated with file descriptor <i>fd</i>. The <i>info</i> pointer is used to return the same information returned by <code>t_open(3NSL)</code>, although not necessarily precisely the same values. This function enables a transport user to access this information during any phase of communication.</p> <p>This argument points to a <code>t_info</code> structure which contains the following members:</p> <pre>t_scalar_t addr;      /*max size in octets of the transport protocol address*/ t_scalar_t options;  /*max number of bytes of protocol-specific options */ t_scalar_t tsdu;     /*max size in octets of a transport service data unit */ t_scalar_t etsdu;    /*max size in octets of an expedited transport service*/                     /*data unit (ETSDU) */ t_scalar_t connect;  /*max number of octets allowed on connection */                     /*establishment functions */ t_scalar_t discon;   /*max number of octets of data allowed on t_snddis() */                     /*and t_rcvdis() functions */ t_scalar_t servtype; /*service type supported by the transport provider */ t_scalar_t flags;    /*other info about the transport provider */</pre> <p>The values of the fields have the following meanings:</p> <p><i>addr</i>            A value greater than zero indicates the maximum size of a transport protocol address and a value of <code>T_INVALID</code> (-2) specifies that the transport provider does not provide user access to transport protocol addresses.</p> <p><i>options</i>        A value greater than zero indicates the maximum number of bytes of protocol-specific options supported by the provider, and a value of <code>T_INVALID</code> (-2) specifies that the transport provider does not support user-settable options.</p> <p><i>tsdu</i>            A value greater than zero specifies the maximum size in octets of a transport service data unit (TSDU); a value of <code>T_NULL</code> (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 <code>T_INFINITE</code> (-1) specifies that there is no limit on the size in octets of a TSDU; and a value of <code>T_INVALID</code> (-2) specifies that the transfer of normal data is not supported by the transport provider.</p>

## t\_getinfo(3NSL)

<i>etsdu</i>	A value greater than zero specifies the maximum size in octets of an expedited transport service data unit (ETSDU); a value of T_NULL (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 T_INFINITE (-1) specifies that there is no limit on the size (in octets) of an ETSDU; and a value of T_INVALID (-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.
<i>connect</i>	A value greater than zero specifies the maximum number of octets that may be associated with connection establishment functions and a value of T_INVALID (-2) specifies that the transport provider does not allow data to be sent with connection establishment functions.
<i>discon</i>	If the T_ORDRELDATA bit in flags is clear, a value greater than zero specifies the maximum number of octets that may be associated with the t_snddis(3NSL) and t_rcvdis(3NSL) functions, and a value of T_INVALID (-2) specifies that the transport provider does not allow data to be sent with the abortive release functions. If the T_ORDRELDATA bit is set in flags, a value greater than zero specifies the maximum number of octets that may be associated with the t_sndreldata(), t_rcvreldata(), t_snddis(3NSL) and t_rcvdis(3NSL) functions.
<i>servtype</i>	This field specifies the service type supported by the transport provider, as described below.
<i>flags</i>	This is a bit field used to specify other information about the communications provider. If the T_ORDRELDATA bit is set, the communications provider supports sending user data with an orderly release. If the T_SENDZERO bit is set in flags, this indicates that 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(3NSL) function 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 may change as a result of protocol option negotiation during connection establishment (the t\_optmgmt(3NSL) call has no effect on the values returned by t\_getinfo()). These values will only change from the values presented to t\_open(3NSL) after the endpoint enters the T\_DATAXFER state.

The *servtype* field of *info* specifies one of the following values on return:

	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, <code>t_open(3NSL)</code> will return <code>T_INVALID (-1)</code> for <code>etsdu</code> , <code>connect</code> and <code>discon</code> .
<b>RETURN VALUES</b>		Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.
<b>VALID STATES</b>		ALL - apart from <code>T_UNINIT</code> .
<b>ERRORS</b>		On failure, <code>t_errno</code> is set to one of the following:
	TBADF	The specified file descriptor does not refer to a transport endpoint.
	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
	TSYSERR	A system error has occurred during execution of this function.
<b>TLI COMPATIBILITY</b>		The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>		The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header:  <pre>#include &lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>		The <code>t_errno</code> value <code>TPROTO</code> can be set by the XTI interface but not by the TLI interface.
<b>The t_info Structure</b>		For TLI, the <code>t_info</code> structure referenced by <i>info</i> lacks the following structure member:  <pre>t_scalar_t flags; /* other info about the transport provider */</pre> This member was added to <code>struct t_info</code> in the XTI interfaces.  When a value of -1 is observed as the return value in various <code>t_info</code> structure members, it signifies that the transport provider can handle an infinite length buffer for a corresponding attribute, such as address data, option data, <code>TSDU</code> (octet size), <code>ETSDU</code> (octet size), connection data, and disconnection data. The corresponding structure members are <code>addr</code> , <code>options</code> , <code>tsdu</code> , <code>estdu</code> , <code>connect</code> , and <code>discon</code> , respectively.
<b>ATTRIBUTES</b>		See <code>attributes(5)</code> for descriptions of the following attributes:

t\_getinfo(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** t\_alloc(3NSL), t\_open(3NSL), t\_optmgmt(3NSL), t\_rcvdis(3NSL),  
t\_snddis(3NSL), attributes(5)

<b>NAME</b>	t_getprotaddr – get the protocol addresses								
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_getprotaddr(int fd, struct t_bind *boundaddr, struct t_bind                  *peeraddr);</pre>								
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>The <code>t_getprotaddr()</code> function returns local and remote protocol addresses currently associated with the transport endpoint specified by <code>fd</code>. In <code>boundaddr</code> and <code>peeraddr</code> the user specifies <code>maxlen</code>, which is the maximum size (in bytes) of the address buffer, and <code>buf</code> which points to the buffer where the address is to be placed. On return, the <code>buf</code> field of <code>boundaddr</code> points to the address, if any, currently bound to <code>fd</code>, and the <code>len</code> field specifies the length of the address. If the transport endpoint is in the <code>T_UNBND</code> state, zero is returned in the <code>len</code> field of <code>boundaddr</code>. The <code>buf</code> field of <code>peeraddr</code> points to the address, if any, currently connected to <code>fd</code>, and the <code>len</code> field specifies the length of the address. If the transport endpoint is not in the <code>T_DATAXFER</code>, <code>T_INREL</code>, <code>T_OUTCON</code> or <code>T_OUTREL</code> states, zero is returned in the <code>len</code> field of <code>peeraddr</code>. If the <code>maxlen</code> field of <code>boundaddr</code> or <code>peeraddr</code> is set to zero, no address is returned.</p>								
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate the error.								
<b>VALID STATES</b>	ALL - apart from <code>T_UNINIT</code> .								
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following:								
	<table border="0"> <tr> <td style="padding-right: 20px;">TBADF</td> <td>The specified file descriptor does not refer to a transport endpoint.</td> </tr> <tr> <td>TBUFOVFLW</td> <td>The number of bytes allocated for an incoming argument (<code>maxlen</code>) is greater than 0 but not sufficient to store the value of that argument.</td> </tr> <tr> <td>TPROTO</td> <td>This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<code>t_errno</code>).</td> </tr> <tr> <td>TSYSERR</td> <td>A system error has occurred during execution of this function.</td> </tr> </table>	TBADF	The specified file descriptor does not refer to a transport endpoint.	TBUFOVFLW	The number of bytes allocated for an incoming argument ( <code>maxlen</code> ) is greater than 0 but not sufficient to store the value of that argument.	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).	TSYSERR	A system error has occurred during execution of this function.
TBADF	The specified file descriptor does not refer to a transport endpoint.								
TBUFOVFLW	The number of bytes allocated for an incoming argument ( <code>maxlen</code> ) is greater than 0 but not sufficient to store the value of that argument.								
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).								
TSYSERR	A system error has occurred during execution of this function.								
<b>TLI COMPATIBILITY ATTRIBUTES</b>	<p>In the TLI interface definition, no counterpart of this routine was defined.</p> <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p>								

t\_getprotaddr(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** t\_bind(3NSL), attributes(5)

<b>NAME</b>	t_getstate – get the current state														
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_getstate(int fd);</pre>														
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>The <code>t_getstate()</code> function returns the current state of the provider associated with the transport endpoint specified by <code>fd</code>.</p>														
<b>RETURN VALUES</b>	<p>State is returned upon successful completion. Otherwise, a value of <code>-1</code> is returned and <code>t_errno</code> is set to indicate an error. The current state is one of the following:</p> <table border="0"> <tr> <td>T_UNBND</td> <td>Unbound.</td> </tr> <tr> <td>T_IDLE</td> <td>Idle.</td> </tr> <tr> <td>T_OUTCON</td> <td>Outgoing connection pending.</td> </tr> <tr> <td>T_INCON</td> <td>Incoming connection pending.</td> </tr> <tr> <td>T_DATAXFER</td> <td>Data transfer.</td> </tr> <tr> <td>T_OUTREL</td> <td>Outgoing direction orderly release sent.</td> </tr> <tr> <td>T_INREL</td> <td>Incoming direction orderly release received.</td> </tr> </table> <p>If the provider is undergoing a state transition when <code>t_getstate()</code> is called, the function will fail.</p>	T_UNBND	Unbound.	T_IDLE	Idle.	T_OUTCON	Outgoing connection pending.	T_INCON	Incoming connection pending.	T_DATAXFER	Data transfer.	T_OUTREL	Outgoing direction orderly release sent.	T_INREL	Incoming direction orderly release received.
T_UNBND	Unbound.														
T_IDLE	Idle.														
T_OUTCON	Outgoing connection pending.														
T_INCON	Incoming connection pending.														
T_DATAXFER	Data transfer.														
T_OUTREL	Outgoing direction orderly release sent.														
T_INREL	Incoming direction orderly release received.														
<b>ERRORS</b>	<p>On failure, <code>t_errno</code> is set to one of the following:</p> <table border="0"> <tr> <td>TBADF</td> <td>The specified file descriptor does not refer to a transport endpoint.</td> </tr> <tr> <td>TPROTO</td> <td>This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<code>t_errno</code>).</td> </tr> <tr> <td>TSTATECHNG</td> <td>The transport provider is undergoing a transient state change.</td> </tr> <tr> <td>TSYSERR</td> <td>A system error has occurred during execution of this function.</td> </tr> </table>	TBADF	The specified file descriptor does not refer to a transport endpoint.	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).	TSTATECHNG	The transport provider is undergoing a transient state change.	TSYSERR	A system error has occurred during execution of this function.						
TBADF	The specified file descriptor does not refer to a transport endpoint.														
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).														
TSTATECHNG	The transport provider is undergoing a transient state change.														
TSYSERR	A system error has occurred during execution of this function.														
<b>TLI COMPATIBILITY</b>	<p>The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.</p>														
<b>Interface Header</b>	<p>The XTI interfaces use the header file, <code>xti.h</code>. TLI interfaces should <i>not</i> use this header. They should use the header:</p> <pre>#include &lt;tiuser.h&gt;</pre>														

t\_getstate(3NSL)

**Error Description  
Values**

The t\_errno value that can be set by the XTI interface and cannot be set by the TLI interface is:

TPROTO

**ATTRIBUTES**

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**

t\_open(3NSL), attributes(5)



<b>NAME</b>	t_listen – listen for a connection indication				
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_listen(int fd, struct t_call *call);</pre>				
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function listens for a connection indication from a calling transport user. The argument <i>fd</i> identifies the local transport endpoint where connection indications arrive, and on return, <i>call</i> contains information describing the connection indication. The parameter <i>call</i> points to a <code>t_call</code> structure which contains the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata; int sequence;</pre> <p>In <i>call</i>, <i>addr</i> returns the protocol address of the calling transport user. This address is in a format usable in future calls to <code>t_connect(3NSL)</code>. Note, however that <code>t_connect(3NSL)</code> may fail for other reasons, for example <code>TADDRBUSY</code>. <i>opt</i> returns options associated with the connection indication, <i>udata</i> returns any user data sent by the caller on the connection request, and <i>sequence</i> is a number that uniquely identifies the returned connection indication. The value of <i>sequence</i> enables the user to listen for multiple connection indications before responding to any of them.</p> <p>Since this function returns values for the <i>addr</i>, <i>opt</i> and <i>udata</i> fields of <i>call</i>, the <i>maxlen</i> field of each must be set before issuing the <code>t_listen()</code> to indicate the maximum size of the buffer for each. If the <i>maxlen</i> field of <i>call</i>→<i>addr</i>, <i>call</i>→<i>opt</i> or <i>call</i>→<i>udata</i> is set to zero, no information is returned for this parameter.</p> <p>By default, <code>t_listen()</code> executes in synchronous mode and waits for a connection indication to arrive before returning to the user. However, if <code>O_NONBLOCK</code> is set via <code>t_open(3NSL)</code> or <code>fcntl(2)</code>, <code>t_listen()</code> executes asynchronously, reducing to a poll for existing connection indications. If none are available, it returns <code>-1</code> and sets <code>t_errno</code> to <code>TNODATA</code>.</p>				
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of <code>-1</code> is returned and <code>t_errno</code> is set to indicate an error.				
<b>VALID STATES</b>	<code>T_IDLE</code> , <code>T_INCON</code>				
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following: <table border="0" style="margin-left: 2em;"> <tr> <td style="padding-right: 1em;"><code>TBADF</code></td> <td>The specified file descriptor does not refer to a transport endpoint.</td> </tr> <tr> <td><code>TBADQLEN</code></td> <td>The argument <i>qlen</i> of the endpoint referenced by <i>fd</i> is zero.</td> </tr> </table>	<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.	<code>TBADQLEN</code>	The argument <i>qlen</i> of the endpoint referenced by <i>fd</i> is zero.
<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.				
<code>TBADQLEN</code>	The argument <i>qlen</i> of the endpoint referenced by <i>fd</i> is zero.				

## t\_listen(3NSL)

TBUFOVFLW	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. The provider's state, as seen by the user, changes to T_INCON, and the connection indication information to be returned in <i>call</i> is discarded. The value of <i>sequence</i> returned can be used to do a t_snddis(3NSL).
TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
TNODATA	O_NONBLOCK was set, but no connection indications had been queued.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (t_errno).
TQFULL	The maximum number of outstanding connection indications has been reached for the endpoint referenced by <i>fd</i> . Note that a subsequent call to t_listen() may block until another incoming connection indication is available. This can only occur if at least one of the outstanding connection indications becomes no longer outstanding, for example through a call to t_accept(3NSL).
TSYSERR	A system error has occurred during execution of this function.
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>	The XTI interfaces use the header file, xti.h. TLI interfaces should <i>not</i> use this header. They should use the header:  <pre>#include &lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>	The t_errno values TPROTO, TBADQLEN, and TQFULL can be set by the XTI interface but not by the TLI interface.  A t_errno value that this routine can return under different circumstances than its XTI counterpart is TBUFOVFLW. It can be returned even when the maxlen field of the corresponding buffer has been set to zero.
<b>Option Buffers</b>	The format of the options in an opt buffer is dictated by the transport provider. Unlike the XTI interface, the TLI interface does not fix the buffer format.

t\_listen(3NSL)

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** fcntl(2), t\_accept(3NSL), t\_alloc(3NSL), t\_bind(3NSL), t\_connect(3NSL), t\_open(3NSL), t\_optmgmt(3NSL), t\_rcvconnect(3NSL), t\_snddis(3NSL), attributes(5)

**WARNINGS** Some transport providers do not differentiate between a connection indication and the connection itself. If this is the case, a successful return of t\_listen() indicates an existing connection.

## t\_look(3NSL)

<b>NAME</b>	t_look – look at the current event on a transport endpoint																		
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_look(int fd);</pre>																		
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function returns the current event on the transport endpoint specified by <i>fd</i>. This function enables a transport provider to notify a transport user of an asynchronous event when the user is calling 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 function to be executed.</p> <p>This function also enables a transport user to poll a transport endpoint periodically for asynchronous events.</p>																		
<b>RETURN VALUES</b>	<p>Upon success, <code>t_look()</code> 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:</p> <table><tr><td>T_LISTEN</td><td>Connection indication received.</td></tr><tr><td>T_CONNECT</td><td>Connect confirmation received.</td></tr><tr><td>T_DATA</td><td>Normal data received.</td></tr><tr><td>T_EXDATA</td><td>Expedited data received.</td></tr><tr><td>T_DISCONNECT</td><td>Disconnection received.</td></tr><tr><td>T_UDERR</td><td>Datagram error indication.</td></tr><tr><td>T_ORDREL</td><td>Orderly release indication.</td></tr><tr><td>T_GODATA</td><td>Flow control restrictions on normal data flow that led to a TFLOW error have been lifted. Normal data may be sent again.</td></tr><tr><td>T_GOEXDATA</td><td>Flow control restrictions on expedited data flow that led to a TFLOW error have been lifted. Expedited data may be sent again.</td></tr></table> <p>On failure, -1 is returned and <code>t_errno</code> is set to indicate the error.</p>	T_LISTEN	Connection indication received.	T_CONNECT	Connect confirmation received.	T_DATA	Normal data received.	T_EXDATA	Expedited data received.	T_DISCONNECT	Disconnection received.	T_UDERR	Datagram error indication.	T_ORDREL	Orderly release indication.	T_GODATA	Flow control restrictions on normal data flow that led to a TFLOW error have been lifted. Normal data may be sent again.	T_GOEXDATA	Flow control restrictions on expedited data flow that led to a TFLOW error have been lifted. Expedited data may be sent again.
T_LISTEN	Connection indication received.																		
T_CONNECT	Connect confirmation received.																		
T_DATA	Normal data received.																		
T_EXDATA	Expedited data received.																		
T_DISCONNECT	Disconnection received.																		
T_UDERR	Datagram error indication.																		
T_ORDREL	Orderly release indication.																		
T_GODATA	Flow control restrictions on normal data flow that led to a TFLOW error have been lifted. Normal data may be sent again.																		
T_GOEXDATA	Flow control restrictions on expedited data flow that led to a TFLOW error have been lifted. Expedited data may be sent again.																		
<b>VALID STATES</b>	ALL - apart from T_UNINIT.																		
<b>ERRORS</b>	<p>On failure, <code>t_errno</code> is set to one of the following:</p> <table><tr><td>TBADF</td><td>The specified file descriptor does not refer to a transport endpoint.</td></tr></table>	TBADF	The specified file descriptor does not refer to a transport endpoint.																
TBADF	The specified file descriptor does not refer to a transport endpoint.																		

t\_look(3NSL)

TPROTO This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (t\_errno).

TSYSERR A system error has occurred during execution of this function.

**TLI COMPATIBILITY**

The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header**

The XTI interfaces use the header file, xti.h. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Return Values**

The return values that are defined by the XTI interface and cannot be returned by the TLI interface are:

T\_GODATA  
T\_GOEXDATA

**Error Description Values**

The t\_errno value that can be set by the XTI interface and cannot be set by the TLI interface is:

TPROTO

**ATTRIBUTES**

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**

t\_open(3NSL), t\_snd(3NSL), t\_sndudata(3NSL), attributes(5)

## t\_open(3NSL)

<b>NAME</b>	t_open – establish a transport endpoint
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt; #include &lt;fcntl.h&gt;  int t_open(const char *name, int oflag, struct t_info *info);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>The <code>t_open()</code> function must be called as the first step in the initialization of a transport endpoint. This function establishes a transport endpoint by supplying a transport provider identifier that indicates a particular transport provider, that is, transport protocol, and returning a file descriptor that identifies that endpoint.</p> <p>The argument <i>name</i> points to a transport provider identifier and <i>oflag</i> identifies any open flags, as in <code>open(2)</code>. The argument <i>oflag</i> is constructed from <code>O_RDWR</code> optionally bitwise inclusive-OR'ed with <code>O_NONBLOCK</code>. These flags are defined by the header <code>&lt;fcntl.h&gt;</code>. The file descriptor returned by <code>t_open()</code> will be used by all subsequent functions to identify the particular local transport endpoint.</p> <p>This function also returns various default characteristics of the underlying transport protocol by setting fields in the <code>t_info</code> structure. This argument points to a <code>t_info</code> which contains the following members:</p> <pre>t_scalar_t addr;          /* max size of the transport protocol address */ t_scalar_t options;       /* max number of bytes of */                           /* protocol-specific options */ t_scalar_t tsdu;          /* max size of a transport service data */                           /* unit (TSDU) */ t_scalar_t etsdu;         /* max size of an expedited transport */                           /* service data unit (ETSDU) */ t_scalar_t connect;       /* max amount of data allowed on */                           /* connection establishment functions */ t_scalar_t discon;        /* max amount of data allowed on */                           /* t_snddis() and t_rcvdis() functions */ t_scalar_t servtype;      /* service type supported by the */                           /* transport provider */ t_scalar_t flags;         /* other info about the transport provider */</pre> <p>The values of the fields have the following meanings:</p> <p><i>addr</i>                    A value greater than zero (<code>T_NULL</code>) indicates the maximum size of a transport protocol address and a value of <code>-2</code> (<code>T_INVALID</code>) specifies that the transport provider does not provide user access to transport protocol addresses.</p> <p><i>options</i>                A value greater than zero (<code>T_NULL</code>) indicates the maximum number of bytes of protocol-specific options supported by the</p>

	provider, and a value of <code>-2</code> ( <code>T_INVALID</code> ) specifies that the transport provider does not support user-settable options.
<i>tsdu</i>	A value greater than zero ( <code>T_NULL</code> ) specifies the maximum size of a transport service data unit (TSDU); a value of zero ( <code>T_NULL</code> ) specifies that the transport provider does not support the concept of TSU, although it does support the sending of a data stream with no logical boundaries preserved across a connection; a value of <code>-1</code> ( <code>T_INFINITE</code> ) specifies that there is no limit to the size of a TSU; and a value of <code>-2</code> ( <code>T_INVALID</code> ) specifies that the transfer of normal data is not supported by the transport provider.
<i>etsdu</i>	A value greater than zero ( <code>T_NULL</code> ) specifies the maximum size of an expedited transport service data unit (ETSDU); a value of zero ( <code>T_NULL</code> ) 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 <code>-1</code> ( <code>T_INFINITE</code> ) specifies that there is no limit on the size of an ETSDU; and a value of <code>-2</code> ( <code>T_INVALID</code> ) 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.
<i>connect</i>	A value greater than zero ( <code>T_NULL</code> ) specifies the maximum amount of data that may be associated with connection establishment functions, and a value of <code>-2</code> ( <code>T_INVALID</code> ) specifies that the transport provider does not allow data to be sent with connection establishment functions.
<i>discon</i>	If the <code>T_ORDRELDATA</code> bit in <code>flags</code> is clear, a value greater than zero ( <code>T_NULL</code> ) specifies the maximum amount of data that may be associated with the <code>t_snddis(3NSL)</code> and <code>t_rcvdis(3NSL)</code> functions, and a value of <code>-2</code> ( <code>T_INVALID</code> ) specifies that the transport provider does not allow data to be sent with the abortive release functions. If the <code>T_ORDRELDATA</code> bit is set in <code>flags</code> , a value greater than zero ( <code>T_NULL</code> ) specifies the maximum number of octets that may be associated with the <code>t_sndreldata()</code> , <code>t_rcvreldata()</code> , <code>t_snddis(3NSL)</code> and <code>t_rcvdis(3NSL)</code> functions.
<i>servtype</i>	This field specifies the service type supported by the transport provider, as described below.
<i>flags</i>	This is a bit field used to specify other information about the communications provider. If the <code>T_ORDRELDATA</code> bit is set, the communications provider supports user data to be sent with an orderly release. If the <code>T_SENDZERO</code> bit is set in <code>flags</code> , this indicates the underlying transport provider supports the sending of zero-length TSUs.

## t\_open(3NSL)

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(3NSL)` function may be used to allocate these buffers. An error will result if a transport user exceeds the allowed data size on any function.

The *servtype* field of *info* specifies one of the following values on return:

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, <code>t_open()</code> will return <code>-2</code> ( <code>T_INVALID</code> ) for <i>etsdu</i> , <i>connect</i> and <i>discon</i> .

A single transport endpoint may support only one of the above services at one time.

If *info* is set to a null pointer by the transport user, no protocol information is returned by `t_open()`.

**RETURN VALUES** A valid file descriptor is returned upon successful completion. Otherwise, a value of `-1` is returned and `t_errno` is set to indicate an error.

**VALID STATES** T\_UNINIT.

**ERRORS** On failure, `t_errno` is set to the following:

TBADFLAG	An invalid flag is specified.
TBADNAME	Invalid transport provider name.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI COMPATIBILITY** The XTI and TLI interface definitions have common names but use different header files. This and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header** The XTI interfaces use the `xti.h` TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Error Description Values** The `t_errno` values `TPROTO` and `TBADNAME` can be set by the XTI interface but cannot be set by the TLI interface.



**Notes** For TLI, the `t_info` structure referenced by *info* lacks the following structure member:

```
t_scalar_t flags; /* other info about the transport provider */
```

This member was added to `struct t_info` in the XTI interfaces.

When a value of `-1` is observed as the return value in various `t_info` structure members, it signifies that the transport provider can handle an infinite length buffer for a corresponding attribute, such as address data, option data, TSDU (octet size), ETSDU (octet size), connection data, and disconnection data. The corresponding structure members are `addr`, `options`, `tsdu`, `estdu`, `connect`, and `discon`, respectively.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `open(2)`, `attributes(5)`

## t\_optmgmt(3NSL)

<b>NAME</b>	t_optmgmt – manage options for a transport endpoint
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_optmgmt(int fd, const struct t_optmgmt *req, struct t_optmgmt               *ret);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>The <code>t_optmgmt()</code> function enables a transport user to retrieve, verify or negotiate protocol options with the transport provider. The argument <code>fd</code> identifies a transport endpoint.</p> <p>The <code>req</code> and <code>ret</code> arguments point to a <code>t_optmgmt</code> structure containing the following members:</p> <pre>struct netbuf opt; t_scalar_t    flags;</pre> <p>The <code>opt</code> field identifies protocol options and the <code>flags</code> field is used to specify the action to take with those options.</p> <p>The options are represented by a <code>netbuf</code> structure in a manner similar to the address in <code>t_bind(3NSL)</code>. The argument <code>req</code> is used to request a specific action of the provider and to send options to the provider. The argument <code>len</code> specifies the number of bytes in the options, <code>buf</code> points to the options buffer, and <code>maxlen</code> has no meaning for the <code>req</code> argument. The transport provider may return options and flag values to the user through <code>ret</code>. For <code>ret</code>, <code>maxlen</code> specifies the maximum size of the options buffer and <code>buf</code> points to the buffer where the options are to be placed. If <code>maxlen</code> in <code>ret</code> is set to zero, no options values are returned. On return, <code>len</code> specifies the number of bytes of options returned. The value in <code>maxlen</code> has no meaning for the <code>req</code> argument, but must be set in the <code>ret</code> argument to specify the maximum number of bytes the options buffer can hold.</p> <p>Each option in the options buffer is of the form <code>struct t_opthdr</code> possibly followed by an option value.</p> <p>The <code>level</code> field of <code>struct t_opthdr</code> identifies the XTI level or a protocol of the transport provider. The <code>name</code> field identifies the option within the level, and <code>len</code> contains its total length; that is, the length of the option header <code>t_opthdr</code> plus the length of the option value. If <code>t_optmgmt()</code> is called with the action <code>T_NEGOTIATE</code> set, the <code>status</code> field of the returned options contains information about the success or failure of a negotiation.</p> <p>Several options can be concatenated. The option user has, however to ensure that each options header and value part starts at a boundary appropriate for the architecture&amp;hyphen;specific alignment rules. The macros <code>T_OPT_FIRSTHDR(nbp)</code>, <code>T_OPT_NEXTHDR(nbp,tohp)</code>, <code>T_OPT_DATA(tohp)</code> are provided for that purpose.</p>

T\_OPT\_DATA (nhp)

If argument is a pointer to a `t_opthdr` structure, this macro returns an unsigned character pointer to the data associated with the `t_opthdr`.

T\_OPT\_NEXTHDR (nbp, tohp)

If the first argument is a pointer to a `netbuf` structure associated with an option buffer and second argument is a pointer to a `t_opthdr` structure within that option buffer, this macro returns a pointer to the next `t_opthdr` structure or a null pointer if this `t_opthdr` is the last `t_opthdr` in the option buffer.

T\_OPT\_FIRSTHDR (tohp)

If the argument is a pointer to a `netbuf` structure associated with an option buffer, this macro returns the pointer to the first `t_opthdr` structure in the associated option buffer, or a null pointer if there is no option buffer associated with this `netbuf` or if it is not possible or the associated option buffer is too small to accommodate even the first aligned option header.

`T_OPT_FIRSTHDR` is useful for finding an appropriately aligned start of the option buffer. `T_OPT_NEXTHDR` is useful for moving to the start of the next appropriately aligned option in the option buffer. Note that `OPT_NEXTHDR` is also available for backward compatibility requirements. `T_OPT_DATA` is useful for finding the start of the data part in the option buffer where the contents of its values start on an appropriately aligned boundary.

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 `t_optmgmt()` request will fail with `TBADOPT`. If the error is detected, some options have possibly been successfully

## t\_optmgmt(3NSL)

### T\_NEGOTIATE

negotiated. The transport user can check the current status by calling `t_optmgmt()` with the `T_CURRENT` flag set.

The *flags* field of *req* must specify one of the following actions:

This action enables the transport user to negotiate option values.

The user specifies the options of interest and their values in the buffer specified by *req*→*opt.buf* and *req*→*opt.len*. The negotiated option values are returned in the buffer pointed to by *ret*→*opt.buf*. The *status* field of each returned option is set to indicate the result of the negotiation. The value is `T_SUCCESS` if the proposed value was negotiated, `T_PARTSUCCESS` if a degraded value was negotiated, `T_FAILURE` if the negotiation failed (according to the negotiation rules), `T_NOTSUPPORT` if the transport provider does not support this option or illegally requests negotiation of a privileged option, and `T_READONLY` if modification of a read-only option was requested. If the status is `T_SUCCESS`, `T_FAILURE`, `T_NOTSUPPORT` or `T_READONLY`, the returned option value is the same as the one requested on input.

The overall result of the negotiation is returned in *ret*→*flags*.

This field contains the worst single result, whereby the rating is done according to the order `T_NOTSUPPORT`, `T_READONLY`, `T_FAILURE`, `T_PARTSUCCESS`, `T_SUCCESS`. The value `T_NOTSUPPORT` is the worst result and `T_SUCCESS` is the best.

For each level, the option `T_ALLOPT` can be requested on input. No value is given with this option; only the `t_opthdr` part is specified. This input requests to negotiate all supported options of this level to their default values. The result is returned option by option in *ret*→*opt.buf*. Note that

t\_optmgmt(3NSL)

depending on the state of the transport endpoint, not all requests to negotiate the default value may be successful.

T\_CHECK

This action enables the user to verify whether the options specified in *req* are supported by the transport provider. If an option is specified with no option value (it consists only of a `t_opthdr` structure), the option is returned with its *status* field set to `T_SUCCESS` if it is supported, `T_NOTSUPPORT` if it is not or needs additional user privileges, and `T_READONLY` if it is read-only (in the current XTI state). No option value is returned.

If an option is specified with an option value, the *status* field of the returned option has the same value, as if the user had tried to negotiate this value with `T_NEGOTIATE`. If the status is `T_SUCCESS`, `T_FAILURE`, `T_NOTSUPPORT` or `T_READONLY`, the returned option value is the same as the one requested on input.

The overall result of the option checks is returned in *ret*→*flags*. This field contains the worst single result of the option checks, whereby the rating is the same as for `T_NEGOTIATE`.

Note that no negotiation takes place. All currently effective option values remain unchanged.

T\_DEFAULT

This action enables the transport user to retrieve the default option values. The user specifies the options of interest in *req*→*opt.buf*. The option values are irrelevant and will be ignored; it is sufficient to specify the `t_opthdr` part of an option only. The default values are then returned in *ret*→*opt.buf*.

The *status* field returned is `T_NOTSUPPORT` if the protocol level does not support this option or the transport user illegally requested a privileged option, `T_READONLY`

t\_optmgmt(3NSL)

T\_CURRENT

if the option is read-only, and set to T\_SUCCESS in all other cases. The overall result of the request is returned in *ret→flags*. This field contains the worst single result, whereby the rating is the same as for T\_NEGOTIATE.

For each level, the option T\_ALLOPT can be requested on input. All supported options of this level with their default values are then returned. In this case, *ret→opt.maxlen* must be given at least the value *info→options* before the call. See *t\_getinfo(3NSL)* and *t\_open(3NSL)*.

This action enables the transport user to retrieve the currently effective option values. The user specifies the options of interest in *req→opt.buf*. The option values are irrelevant and will be ignored; it is sufficient to specify the *t\_opthdr* part of an option only. The currently effective values are then returned in *req→opt.buf*.

The *status* field returned is T\_NOTSUPPORT if the protocol level does not support this option or the transport user illegally requested a privileged option, T\_READONLY if the option is read-only, and set to T\_SUCCESS in all other cases. The overall result of the request is returned in *ret→flags*. This field contains the worst single result, whereby the rating is the same as for T\_NEGOTIATE.

For each level, the option T\_ALLOPT can be requested on input. All supported options of this level with their currently effective values are then returned.

The option T\_ALLOPT can only be used with *t\_optmgmt()* and the actions T\_NEGOTIATE, T\_DEFAULT and T\_CURRENT. It can be used with any supported level and addresses all supported options of this level. The option has no value; it consists of a *t\_opthdr* only. Since in a *t\_optmgmt()* call only

## t\_optmgmt(3NSL)

options of one level may be addressed, this option should not be requested together with other options. The function returns as soon as this option has been processed.

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.

Transport providers may not be able to provide an interface capable of supporting T\_NEGOTIATE and/or T\_CHECK functionalities. When this is the case, the error TNOTSUPPORT is returned.

The function `t_optmgmt()` may block under various circumstances and depending on the implementation. The function will block, for instance, if the protocol addressed by the call resides on a separate controller. It may also block due to flow control constraints; that is, if data sent previously across this transport endpoint has not yet been fully processed. If the function is interrupted by a signal, the option negotiations that have been done so far may remain valid. The behavior of the function is not changed if `O_NONBLOCK` is set.

**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.

**VALID STATES** ALL - apart from T\_UNINIT.

**ERRORS** On failure, `t_errno` is set to one of the following:

TBADF	The specified file descriptor does not refer to a transport endpoint.
TBADFLAG	An invalid flag was specified.
TBADOPT	The specified options were in an incorrect format or contained illegal information.
TBUFOVFLW	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.
TNOTSUPPORT	This action is not supported by the transport provider.

## t\_optmgmt(3NSL)

	TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
	TSYSERR	A system error has occurred during execution of this function.
<b>TLI COMPATIBILITY</b>		The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>		The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header:  <pre>#include &lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>		The <code>t_errno</code> value TPROTO can be set by the XTI interface but not by the TLI interface.  The <code>t_errno</code> values that this routine can return under different circumstances than its XTI counterpart are TACCES and TBUFOVFLW.  TACCES can be returned to indicate that the user does not have permission to negotiate the specified options.  TBUFOVFLW can be returned even when the <code>maxlen</code> field of the corresponding buffer has been set to zero.
<b>Option Buffers</b>		The format of the options in an <code>opt</code> buffer is dictated by the transport provider. Unlike the XTI interface, the TLI interface does not fix the buffer format. The macros <code>T_OPT_DATA</code> , <code>T_OPT_NEXTHDR</code> , and <code>T_OPT_FIRSTHDR</code> described for XTI are not available for use by TLI interfaces.
<b>Actions</b>		The semantic meaning of various action values for the <code>flags</code> field of <i>req</i> differs between the TLI and XTI interfaces. TLI interface users should heed the following descriptions of the actions:  T_NEGOTIATE This action enables the user to negotiate the values of the options specified in <i>req</i> with the transport provider. The provider will evaluate the requested options and negotiate the values, returning the negotiated values through <i>ret</i> .  T_CHECK This action enables the user to verify whether the options specified in <i>req</i> are supported by the transport provider. On return, the <code>flags</code> field of <i>ret</i> will have either <code>T_SUCCESS</code> or <code>T_FAILURE</code> set to indicate to the user whether the options are supported. These flags are only meaningful for the <code>T_CHECK</code> request.



t\_optmgmt(3NSL)

T\_DEFAULT This action enables a user to retrieve the default options supported by the transport provider into the *opt* field of *ret*. In *req*, the *len* field of *opt* must be zero and the *buf* field may be NULL.

**Connectionless Mode**

If issued as part of the connectionless mode service, t\_optmgmt () may block due to flow control constraints. The function will not complete until the transport provider has processed all previously sent data units.

**ATTRIBUTES**

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**

close(2), poll(2), select(3C), t\_accept(3NSL), t\_alloc(3NSL), t\_bind(3NSL), t\_close(3NSL), t\_connect(3NSL), t\_getinfo(3NSL), t\_listen(3NSL), t\_open(3NSL), t\_rcv(3NSL), t\_rcvconnect(3NSL), t\_rcvudata(3NSL), t\_snddis(3NSL), attributes(5)

## t\_rcv(3NSL)

<b>NAME</b>	t_rcv – receive data or expedited data sent over a connection
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcv(int fd, void *buf, unsigned int nbytes, int *flags);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function receives either normal or expedited data. The argument <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, and <i>nbytes</i> specifies the size of the receive buffer. The argument <i>flags</i> may be set on return from <code>t_rcv()</code> and specifies optional flags as described below.</p> <p>By default, <code>t_rcv()</code> operates in synchronous mode and will wait for data to arrive if none is currently available. However, if <code>O_NONBLOCK</code> is set by means of <code>t_open(3NSL)</code> or <code>fcntl(2)</code>, <code>t_rcv()</code> will execute in asynchronous mode and will fail if no data is available. See <code>TNODATA</code> below.</p> <p>On return from the call, if <code>T_MORE</code> is set in <i>flags</i>, this 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 <code>t_rcv()</code> calls. In the asynchronous mode, or under unusual conditions (for example, the arrival of a signal or <code>T_EXDATA</code> event), the <code>T_MORE</code> flag may be set on return from the <code>t_rcv()</code> call even when the number of bytes received is less than the size of the receive buffer specified. Each <code>t_rcv()</code> with the <code>T_MORE</code> flag set indicates that another <code>t_rcv()</code> must follow to get more data for the current TSDU. The end of the TSDU is identified by the return of a <code>t_rcv()</code> call with the <code>T_MORE</code> flag not set. If the transport provider does not support the concept of a TSDU as indicated in the <i>info</i> argument on return from <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>, the <code>T_MORE</code> flag is not meaningful and should be ignored. If <i>nbytes</i> is greater than zero on the call to <code>t_rcv()</code>, <code>t_rcv()</code> will return 0 only if the end of a TSDU is being returned to the user.</p> <p>On return, the data is expedited if <code>T_EXPEDITED</code> is set in <i>flags</i>. If <code>T_MORE</code> is also set, it indicates that the number of expedited bytes exceeded <i>nbytes</i>, a signal has interrupted the call, or that an entire ETSDU was not available (only for transport protocols that support fragmentation of ETSDUs). The rest of the ETSDU will be returned by subsequent calls to <code>t_rcv()</code> which will return with <code>T_EXPEDITED</code> set in <i>flags</i>. The end of the ETSDU is identified by the return of a <code>t_rcv()</code> call with <code>T_EXPEDITED</code> set and <code>T_MORE</code> cleared. If the entire ETSDU is not available it is possible for normal data fragments to be returned between the initial and final fragments of an ETSDU.</p>

t\_rcv(3NSL)

If a signal arrives, t\_rcv() returns, giving the user any data currently available. If no data is available, t\_rcv() returns -1, sets t\_errno to TSYSERR and errno to EINTR. If some data is available, t\_rcv() returns the number of bytes received and T\_MORE is set in flags.

In synchronous mode, the only way for the user to be notified of the arrival of normal or expedited data is to issue this function or check for the T\_DATA or T\_EXDATA events using the t\_look(3NSL) function. Additionally, the process can arrange to be notified by means of the EM interface.

**RETURN VALUES** On successful completion, t\_rcv() returns the number of bytes received. Otherwise, it returns 1 on failure and t\_errno is set to indicate the error.

**VALID STATES** T\_DATAXFER, T\_OUTREL.

**ERRORS** On failure, t\_errno is set to one of the following:

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.
TNODATA	O_NONBLOCK was set, but no data is currently available from the transport provider.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (t_errno).
TSYSERR	A system error has occurred during execution of this function.

**TLI COMPATIBILITY** The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header** The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Error Description Values** The t\_errno value that can be set by the XTI interface and cannot be set by the TLI interface is:

TPROTO

t\_rcv(3NSL)

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `fcntl(2)`, `t_getinfo(3NSL)`, `t_look(3NSL)`, `t_open(3NSL)`, `t_snd(3NSL)`, `attributes(5)`

<b>NAME</b>	t_rcvconnect – receive the confirmation from a connection request
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcvconnect(int fd, struct t_call *call);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function enables a calling transport user to determine the status of a previously sent connection request and is used in conjunction with <code>t_connect(3NSL)</code> to establish a connection in asynchronous mode, and to complete a synchronous <code>t_connect(3NSL)</code> call that was interrupted by a signal. The connection will be established on successful completion of this function.</p> <p>The argument <i>fd</i> identifies the local transport endpoint where communication will be established, and <i>call</i> contains information associated with the newly established connection. The argument <i>call</i> points to a <code>t_call</code> structure which contains the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata; int sequence;</pre> <p>In <i>call</i>, <i>addr</i> returns the protocol address associated with the responding transport endpoint, <i>opt</i> presents any options associated with the connection, <i>udata</i> points to optional user data that may be returned by the destination transport user during connection establishment, and <i>sequence</i> has no meaning for this function.</p> <p>The <i>maxlen</i> field of each argument must be set before issuing this function to indicate the maximum size of the buffer for each. However, <i>maxlen</i> can be set to zero, in which case no information to this specific argument is given to the user on the return from <code>t_rcvconnect()</code>. If <i>call</i> is set to <code>NULL</code>, no information at all is returned. By default, <code>t_rcvconnect()</code> executes in synchronous mode and waits for the connection to be established before returning. On return, the <i>addr</i>, <i>opt</i> and <i>udata</i> fields reflect values associated with the connection.</p> <p>If <code>O_NONBLOCK</code> is set by means of <code>t_open(3NSL)</code> or <code>fcntl(2)</code>, <code>t_rcvconnect()</code> executes in asynchronous mode, and reduces to a poll for existing connection confirmations. If none are available, <code>t_rcvconnect()</code> fails and returns immediately without waiting for the connection to be established. See <code>TNODATA</code> below. In this case, <code>t_rcvconnect()</code> must be called again to complete the connection establishment phase and retrieve the information returned in <i>call</i>.</p>
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.
<b>VALID STATES</b>	<code>T_OUTCON</code> .

t\_rcvconnect(3NSL)

**ERRORS**

On failure, `t_errno` is set to one of the following:

TBADF	The specified file descriptor does not refer to a transport endpoint.
TBUFOVFLW	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 connection information to be returned in <i>call</i> will be discarded. The provider's state, as seen by the user, will be changed to <code>T_DATAXFER</code> .
TLOOK	An asynchronous event has occurred on this transport connection and requires immediate attention.
TNODATA	<code>O_NONBLOCK</code> was set, but a connection confirmation has not yet arrived.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI COMPATIBILITY**

The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header**

The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include<tiuser.h>
```

**Error Description Values**

The `t_errno` value `TPROTO` can be set by the XTI interface but not by the TLI interface.

A `t_errno` value that this routine can return under different circumstances than its XTI counterpart is `TBUFOVFLW`. It can be returned even when the `maxlen` field of the corresponding buffer has been set to zero.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

t\_rcvconnect(3NSL)

**SEE ALSO** fcntl(2), t\_accept(3NSL), t\_alloc(3NSL), t\_bind(3NSL), t\_connect(3NSL),  
t\_listen(3NSL), t\_open(3NSL), t\_optmgmt(3NSL), attributes(5)

## t\_rcvdis(3NSL)

<b>NAME</b>	t_rcvdis – retrieve information from disconnection				
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcvdis(int fd, struct t_discon *discon);</pre>				
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function is used to identify the cause of a disconnection and to retrieve any user data sent with the disconnection. The argument <i>fd</i> identifies the local transport endpoint where the connection existed, and <i>discon</i> points to a <code>t_discon</code> structure containing the following members:</p> <pre>struct netbuf udata; int reason; int sequence;</pre> <p>The field <i>reason</i> specifies the reason for the disconnection through a protocol-dependent reason code, <i>udata</i> identifies any user data that was sent with the disconnection, and <i>sequence</i> may identify an outstanding connection indication with which the disconnection is associated. The field <i>sequence</i> is only meaningful when <code>t_rcvdis()</code> is issued by a passive transport user who has executed one or more <code>t_listen(3NSL)</code> functions and is processing the resulting connection indications. If a disconnection indication occurs, <i>sequence</i> can be used to identify which of the outstanding connection indications is associated with the disconnection.</p> <p>The <i>maxlen</i> field of <i>udata</i> may be set to zero, if the user does not care about incoming data. If, in addition, the user does not need to know the value of <i>reason</i> or <i>sequence</i>, <i>discon</i> may be set to NULL and any user data associated with the disconnection indication shall be discarded. However, if a user has retrieved more than one outstanding connection indication by means of <code>t_listen(3NSL)</code>, and <i>discon</i> is a null pointer, the user will be unable to identify with which connection indication the disconnection is associated.</p>				
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.				
<b>VALID STATES</b>	T_DATAXFER, T_OUTCON, T_OUTREL, T_INREL, T_INCON( <i>ocnt</i> > 0).				
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following:				
	<table><tr><td>TBADF</td><td>The specified file descriptor does not refer to a transport endpoint.</td></tr><tr><td>TBUFOVFLW</td><td>The number of bytes allocated for incoming data (<i>maxlen</i>) is greater than 0 but not sufficient to store the data. If <i>fd</i> is a passive endpoint with <i>ocnt</i> &gt; 1, it remains in state T_INCON; otherwise, the endpoint state is set to T_IDLE.</td></tr></table>	TBADF	The specified file descriptor does not refer to a transport endpoint.	TBUFOVFLW	The number of bytes allocated for incoming data ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the data. If <i>fd</i> is a passive endpoint with <i>ocnt</i> > 1, it remains in state T_INCON; otherwise, the endpoint state is set to T_IDLE.
TBADF	The specified file descriptor does not refer to a transport endpoint.				
TBUFOVFLW	The number of bytes allocated for incoming data ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the data. If <i>fd</i> is a passive endpoint with <i>ocnt</i> > 1, it remains in state T_INCON; otherwise, the endpoint state is set to T_IDLE.				



TNODIS	No disconnection indication currently exists on the specified transport endpoint.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI COMPATIBILITY**

The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header**

The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Error Description Values**

The `t_errno` values TPROTO and TOUTSTATE can be set by the XTI interface but not by the TLI interface.

A failure return, and a `t_errno` value that this routine can set under different circumstances than its XTI counterpart is TBUFOVFLW. It can be returned even when the `maxlen` field of the corresponding buffer has been set to zero.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**

`t_alloc(3NSL)`, `t_connect(3NSL)`, `t_listen(3NSL)`, `t_open(3NSL)`, `t_snddis(3NSL)`, `attributes(5)`

## t\_rcvrel(3NSL)

<b>NAME</b>	t_rcvrel – acknowledge receipt of an orderly release indication														
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcvrel(int fd);</pre>														
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function is used to receive an orderly release indication for the incoming direction of data transfer. The argument <i>fd</i> identifies the local transport endpoint where the connection exists. After receipt of this indication, the user may not attempt to receive more data by means of <code>t_rcv(3NSL)</code> or <code>t_rcvv()</code>. Such an attempt will fail with <i>t_error</i> set to TOUTSTATE. However, the user may continue to send data over the connection if <code>t_sndrel(3NSL)</code> 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 service type <code>T_COTS_ORD</code> on <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>. Any user data that may be associated with the orderly release indication is discarded when <code>t_rcvrel()</code> is called.</p>														
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.														
<b>VALID STATES</b>	T_DATAXFER, T_OUTREL.														
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following:														
	<table><tr><td>TBADF</td><td>The specified file descriptor does not refer to a transport endpoint.</td></tr><tr><td>TLOOK</td><td>An asynchronous event has occurred on this transport endpoint and requires immediate attention.</td></tr><tr><td>TNOREL</td><td>No orderly release indication currently exists on the specified transport endpoint.</td></tr><tr><td>TNOTSUPPORT</td><td>This function is not supported by the underlying transport provider.</td></tr><tr><td>TOUTSTATE</td><td>The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.</td></tr><tr><td>TPROTO</td><td>This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<code>t_errno</code>).</td></tr><tr><td>TSYSERR</td><td>A system error has occurred during execution of this function.</td></tr></table>	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 function is not supported by the underlying transport provider.	TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).	TSYSERR	A system error has occurred during execution of this function.
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 function is not supported by the underlying transport provider.														
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.														
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).														
TSYSERR	A system error has occurred during execution of this function.														
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.														

t\_rcvrel(3NSL)

**Interface Header** The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include<tiuser.h>
```

**Error Description Values** The `t_errno` values that can be set by the XTI interface and cannot be set by the TLI interface are:

```
TPROTO  
TOUTSTATE
```

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `t_getinfo(3NSL)`, `t_open(3NSL)`, `t_sndrel(3NSL)`, `attributes(5)`

## t\_rcvreldata(3NSL)

<b>NAME</b>	t_rcvreldata – receive an orderly release indication or confirmation containing user data				
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcvreldata(int fd, struct t_discon *discon);</pre>				
<b>DESCRIPTION</b>	<p>This function is used to receive an orderly release indication for the incoming direction of data transfer and to retrieve any user data sent with the release. The argument <i>fd</i> identifies the local transport endpoint where the connection exists, and <i>discon</i> points to a <i>t_discon</i> structure containing the following members:</p> <pre>struct netbuf udata; int reason; int sequence;</pre> <p>After receipt of this indication, the user may not attempt to receive more data by means of <i>t_rcv</i>(3NSL) or <i>t_rcvv</i>(3NSL). Such an attempt will fail with <i>t_error</i> set to TOUTSTATE. However, the user may continue to send data over the connection if <i>t_sndrel</i>(3NSL) or <i>t_sndreldata</i> (3N) has not been called by the user.</p> <p>The field <i>reason</i> specifies the reason for the disconnection through a protocol-dependent <i>reason code</i>, and <i>udata</i> identifies any user data that was sent with the disconnection; the field <i>sequence</i> is not used.</p> <p>If a user does not care if there is incoming data and does not need to know the value of <i>reason</i>, <i>discon</i> may be a null pointer, and any user data associated with the disconnection will be discarded.</p> <p>If <i>discon</i>→<i>udata.maxlen</i> is greater than zero and less than the length of the value, <i>t_rcvreldata</i> () fails with <i>t_errno</i> set to TBUFOVFLW.</p> <p>This function is an optional service of the transport provider, only supported by providers of service type T_COTS_ORD. The flag T_ORDRELDATA in the <i>info</i>→<i>flag</i> field returned by <i>t_open</i>(3NSL) or <i>t_getinfo</i>(3NSL) indicates that the provider supports orderly release user data; when the flag is not set, this function behaves like <i>t_rcvrel</i>(3NSL) and no user data is returned.</p> <p>This function may not be available on all systems.</p>				
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <i>t_errno</i> is set to indicate an error.				
<b>VALID STATES</b>	T_DATAXFER, T_OUTREL.				
<b>ERRORS</b>	On failure, <i>t_errno</i> is set to one of the following:				
	<table><tr><td>TBADF</td><td>The specified file descriptor does not refer to a transport endpoint.</td></tr><tr><td>TBUFOVFLW</td><td>The number of bytes allocated for incoming data (<i>maxlen</i>) is greater than 0 but not sufficient to store the data, and the disconnection information to be returned in <i>discon</i> will be</td></tr></table>	TBADF	The specified file descriptor does not refer to a transport endpoint.	TBUFOVFLW	The number of bytes allocated for incoming data ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the data, and the disconnection information to be returned in <i>discon</i> will be
TBADF	The specified file descriptor does not refer to a transport endpoint.				
TBUFOVFLW	The number of bytes allocated for incoming data ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the data, and the disconnection information to be returned in <i>discon</i> will be				

t\_rcvreldata(3NSL)

discarded. The provider state, as seen by the user, will be changed as if the data was successfully retrieved.

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	Orderly release is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <i>t_errno</i> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI  
COMPATIBILITY  
ATTRIBUTES**

In the TLI interface definition, no counterpart of this routine was defined.

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `t_getinfo(3NSL)`, `t_open(3NSL)`, `t_sndreldata(3NSL)`, `t_rcvrel(3NSL)`, `t_sndrel(3NSL)`, `attributes(5)`

**NOTES** The interfaces `t_sndreldata(3NSL)` and `t_rcvreldata()` are only for use with a specific transport called "minimal OSI," which is not available on the Solaris platform. These interfaces are not available for use in conjunction with Internet Transports (TCP or UDP).

## t\_rcvudata(3NSL)

<b>NAME</b>	t_rcvudata – receive a data unit
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcvudata(int fd, struct t_unitdata *unitdata, int *flags);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function is used in connectionless-mode to receive a data unit from another transport user. The argument <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, and <i>flags</i> is set on return to indicate that the complete data unit was not received. The argument <i>unitdata</i> points to a <code>t_unitdata</code> structure containing the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata;</pre> <p>The <i>maxlen</i> field of <i>addr</i>, <i>opt</i> and <i>udata</i> must be set before calling this function to indicate the maximum size of the buffer for each. If the <i>maxlen</i> field of <i>addr</i> or <i>opt</i> is set to zero, no information is returned in the <i>buf</i> field of this parameter.</p> <p>On return from this call, <i>addr</i> specifies the protocol address of the sending user, <i>opt</i> identifies options that were associated with this data unit, and <i>udata</i> specifies the user data that was received.</p> <p>By default, <code>t_rcvudata()</code> operates in synchronous mode and will wait for a data unit to arrive if none is currently available. However, if <code>O_NONBLOCK</code> is set by means of <code>t_open(3NSL)</code> or <code>fcntl(2)</code>, <code>t_rcvudata()</code> will execute in asynchronous mode and will fail if no data units are available.</p> <p>If the buffer defined in the <i>udata</i> field of <i>unitdata</i> is not large enough to hold the current data unit, the buffer will be filled and <code>T_MORE</code> will be set in <i>flags</i> on return to indicate that another <code>t_rcvudata()</code> should be called to retrieve the rest of the data unit. Subsequent calls to <code>t_rcvudata()</code> will return zero for the length of the address and options until the full data unit has been received.</p> <p>If the call is interrupted, <code>t_rcvudata()</code> will return <code>EINTR</code> and no datagrams will have been removed from the endpoint.</p>
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.
<b>VALID STATES</b>	<code>T_IDLE</code> .
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following:

	TBADF	The specified file descriptor does not refer to a transport endpoint.
	TBUFOVFLW	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 <i>unitdata</i> will be discarded.
	TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
	TNODATA	O_NONBLOCK was set, but no data units are currently available from the transport provider.
	TNOTSUPPORT	This function is not supported by the underlying transport provider.
	TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <i>t_errno</i> ).
	TSYSERR	A system error has occurred during execution of this function.
<b>TLI COMPATIBILITY</b>		The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>		The XTI interfaces use the header file, <i>xti.h</i> . TLI interfaces should <i>not</i> use this header. They should use the header:  <pre>#include&lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>		The <i>t_errno</i> values that can be set by the XTI interface and cannot be set by the TLI interface are:  TPROTO TOUTSTATE  A <i>t_errno</i> value that this routine can return under different circumstances than its XTI counterpart is TBUFOVFLW. It can be returned even when the <i>maxlen</i> field of the corresponding buffer has been set to zero.
<b>Option Buffers</b>		The format of the options in an <i>opt</i> buffer is dictated by the transport provider. Unlike the XTI interface, the TLI interface does not fix the buffer format.
<b>ATTRIBUTES</b>		See <i>attributes(5)</i> for descriptions of the following attributes:

t\_rcvudata(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** fcntl(2), t\_alloc(3NSL), t\_open(3NSL), t\_rcvuderr(3NSL),  
t\_sndudata(3NSL), attributes(5)



<b>NAME</b>	t_rcvuderr – receive a unit data error indication								
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcvuderr(int fd, struct t_uderr *uderr);</pre>								
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function 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. The argument <i>fd</i> identifies the local transport endpoint through which the error report will be received, and <i>uderr</i> points to a <code>t_uderr</code> structure containing the following members:</p> <pre>struct netbuf addr; struct netbuf opt; t_scalar_t error;</pre> <p>The <i>maxlen</i> field of <i>addr</i> and <i>opt</i> must be set before calling this function to indicate the maximum size of the buffer for each. If this field is set to zero for <i>addr</i> or <i>opt</i>, no information is returned in the <i>buf</i> field of this parameter.</p> <p>On return from this call, the <i>addr</i> structure specifies the destination protocol address of the erroneous data unit, the <i>opt</i> structure identifies options that were associated with the data unit, and <i>error</i> specifies a protocol-dependent error code.</p> <p>If the user does not care to identify the data unit that produced an error, <i>uderr</i> may be set to a null pointer, and <code>t_rcvuderr()</code> will simply clear the error indication without reporting any information to the user.</p>								
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.								
<b>VALID STATES</b>	T_IDLE.								
<b>ERRORS</b>	<p>On failure, <code>t_errno</code> is set to one of the following:</p> <table border="0"> <tr> <td style="padding-right: 20px;">TBADF</td> <td>The specified file descriptor does not refer to a transport endpoint.</td> </tr> <tr> <td>TBUFOVFLW</td> <td>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 error information to be returned in <i>uderr</i> will be discarded.</td> </tr> <tr> <td>TNOTSUPPORT</td> <td>This function is not supported by the underlying transport provider.</td> </tr> <tr> <td>TNOUDERR</td> <td>No unit data error indication currently exists on the specified transport endpoint.</td> </tr> </table>	TBADF	The specified file descriptor does not refer to a transport endpoint.	TBUFOVFLW	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 error information to be returned in <i>uderr</i> will be discarded.	TNOTSUPPORT	This function is not supported by the underlying transport provider.	TNOUDERR	No unit data error indication currently exists on the specified transport endpoint.
TBADF	The specified file descriptor does not refer to a transport endpoint.								
TBUFOVFLW	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 error information to be returned in <i>uderr</i> will be discarded.								
TNOTSUPPORT	This function is not supported by the underlying transport provider.								
TNOUDERR	No unit data error indication currently exists on the specified transport endpoint.								

## t\_rcvuderr(3NSL)

TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TSYSERR	A system error has occurred during execution of this function.

### TLI COMPATIBILITY

The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

### Interface Header

The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

### Error Description Values

The `t_errno` values TPROTO and TOUTSTATE can be set by the XTI interface but not by the TLI interface.

A `t_errno` value that this routine can return under different circumstances than its XTI counterpart is TBUFOVFLW. It can be returned even when the `maxLen` field of the corresponding buffer has been set to zero.

### Option Buffers

The format of the options in an `opt` buffer is dictated by the transport provider. Unlike the XTI interface, the TLI interface does not fix the buffer format.

### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

### SEE ALSO

`t_rcvudata(3NSL)`, `t_sndudata(3NSL)`, `attributes(5)`

<b>NAME</b>	t_rcvv – receive data or expedited data sent over a connection and put the data into one or more non-contiguous buffers
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcvv(int fd, struct t_iovec *iov, unsigned int iovcount, int            *flags);</pre>
<b>DESCRIPTION</b>	<p>This function receives either normal or expedited data. The argument <i>fd</i> identifies the local transport endpoint through which data will arrive, <i>iov</i> points to an array of buffer address/buffer size pairs (<i>iov_base</i>, <i>iov_len</i>). The <code>t_rcvv()</code> function receives data into the buffers specified by <i>iov0.iov_base</i>, <i>iov1.iov_base</i>, through <i>iov[iovcount-1].iov_base</i>, always filling one buffer before proceeding to the next.</p> <p>Note that the limit on the total number of bytes available in all buffers passed: <math>iov(0).iov\_len + \dots + iov(iovcount-1).iov\_len</math> may be constrained by implementation limits. If no other constraint applies, it will be limited by <code>INT_MAX</code>. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.</p> <p>The argument <i>iovcount</i> contains the number of buffers which is limited to <code>T_IOV_MAX</code>, which is an implementation-defined value of at least 16. If the limit is exceeded, the function will fail with <code>TBADDATA</code>.</p> <p>The argument <i>flags</i> may be set on return from <code>t_rcvv()</code> and specifies optional flags as described below.</p> <p>By default, <code>t_rcvv()</code> operates in synchronous mode and will wait for data to arrive if none is currently available. However, if <code>O_NONBLOCK</code> is set by means of <code>t_open(3NSL)</code> or <code>fcntl(2)</code>, <code>t_rcvv()</code> will execute in asynchronous mode and will fail if no data is available. See <code>TNODATA</code> below.</p> <p>On return from the call, if <code>T_MORE</code> is set in <i>flags</i>, this 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 <code>t_rcvv()</code> or <code>t_rcv(3NSL)</code> calls. In the asynchronous mode, or under unusual conditions (for example, the arrival of a signal or <code>T_EXDATA</code> event), the <code>T_MORE</code> flag may be set on return from the <code>t_rcvv()</code> call even when the number of bytes received is less than the total size of all the receive buffers. Each <code>t_rcvv()</code> with the <code>T_MORE</code> flag set indicates that another <code>t_rcvv()</code> must follow to get more data for the current TSDU. The end of the TSDU is identified by the return of a <code>t_rcvv()</code> call with the <code>T_MORE</code> flag not set. If the transport provider does not support the concept of a TSDU as indicated in the <i>info</i> argument on return from <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>, the <code>T_MORE</code> flag is not meaningful and should be ignored. If the amount of buffer space passed in <i>iov</i> is greater than zero on the call to <code>t_rcvv()</code>, then <code>t_rcvv()</code> will return 0 only if the end of a TSDU is being returned to the user.</p>

## t\_rcvv(3NSL)

On return, the data is expedited if `T_EXPEDITED` is set in flags. If `T_MORE` is also set, it indicates that the number of expedited bytes exceeded `nbytes`, a signal has interrupted the call, or that an entire ETSDU was not available (only for transport protocols that support fragmentation of ETSDUs). The rest of the ETSDU will be returned by subsequent calls to `t_rcvv()` which will return with `T_EXPEDITED` set in flags. The end of the ETSDU is identified by the return of a `t_rcvv()` call with `T_EXPEDITED` set and `T_MORE` cleared. If the entire ETSDU is not available it is possible for normal data fragments to be returned between the initial and final fragments of an ETSDU.

If a signal arrives, `t_rcvv()` returns, giving the user any data currently available. If no data is available, `t_rcvv()` returns `-1`, sets `t_errno` to `TSYSERR` and `errno` to `EINTR`. If some data is available, `t_rcvv()` returns the number of bytes received and `T_MORE` is set in flags.

In synchronous mode, the only way for the user to be notified of the arrival of normal or expedited data is to issue this function or check for the `T_DATA` or `T_EXDATA` events using the `t_look(3NSL)` function. Additionally, the process can arrange to be notified via the EM interface.

**RETURN VALUES** On successful completion, `t_rcvv()` returns the number of bytes received. Otherwise, it returns `-1` on failure and `t_errno` is set to indicate the error.

**VALID STATES** `T_DATAXFER`, `T_OUTREL`.

**ERRORS** On failure, `t_errno` is set to one of the following:

<code>TBADDATA</code>	<code>iovcount</code> is greater than <code>T_IOV_MAX</code> .
<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.
<code>TLOOK</code>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<code>TNODATA</code>	<code>O_NONBLOCK</code> was set, but no data is currently available from the transport provider.
<code>TNOTSUPPORT</code>	This function is not supported by the underlying transport provider.
<code>TOUTSTATE</code>	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
<code>TPROTO</code>	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
<code>TSYSERR</code>	A system error has occurred during execution of this function.

**TLI COMPATIBILITY ATTRIBUTES** In the TLI interface definition, no counterpart of this routine was defined. See `attributes(5)` for descriptions of the following attributes:

t\_rcvv(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** fcntl(2), t\_getinfo(3NSL), t\_look(3NSL), t\_open(3NSL), t\_rcv(3NSL), t\_snd(3NSL), t\_sndv(3NSL), attributes(5)

## t\_rcvvudata(3NSL)

<b>NAME</b>	t_rcvvudata – receive a data unit into one or more noncontiguous buffers
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_rcvvudata(int fd, struct t_unitdata *unitdata, struct t_iovec                 *iov, unsigned int iovcount, int *flags);</pre>
<b>DESCRIPTION</b>	<p>This function is used in connectionless mode to receive a data unit from another transport user. The argument <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, <i>iovcount</i> contains the number of non-contiguous udata buffers which is limited to T_IOV_MAX, which is an implementation-defined value of at least 16, and <i>flags</i> is set on return to indicate that the complete data unit was not received. If the limit on <i>iovcount</i> is exceeded, the function fails with TBADDDATA. The argument <i>unitdata</i> points to a t_unitdata structure containing the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata; The maxlen field of <i>addr</i> and <i>opt</i> must be set before calling this function to indicate the maximum size of the buffer for each. The <i>udata</i> field of t_unitdata is not used. The <i>iov_len</i> and <i>iov_base</i> fields of "<i>iov0</i>" through <i>iov</i> [<i>iovcount-1</i>] must be set before calling t_rcvvudata() to define the buffer where the userdata will be placed. If the maxlen field of <i>addr</i> or <i>opt</i> is set to zero then no information is returned in the <i>buf</i> field for this parameter.</pre> <p>On return from this call, <i>addr</i> specifies the protocol address of the sending user, <i>opt</i> identifies options that were associated with this data unit, and <i>iov</i>[0].<i>iov_base</i> through <i>iov</i>[<i>iovcount-1</i>].<i>iov_base</i> contains the user data that was received. The return value of t_rcvvudata() is the number of bytes of user data given to the user.</p> <p>Note that the limit on the total number of bytes available in all buffers passed:</p> <pre>iov(0).iov_len + .. + iov(iovcount-1).iov_len</pre> may be constrained by implementation limits. If no other constraint applies, it will be limited by INT_MAX. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions. <p>By default, t_rcvvudata() operates in synchronous mode and waits for a data unit to arrive if none is currently available. However, if O_NONBLOCK is set by means of t_open(3NSL) orfcntl(2), t_rcvvudata() executes in asynchronous mode and fails if no data units are available.</p> <p>If the buffers defined in the <i>iov</i>[] array are not large enough to hold the current data unit, the buffers will be filled and T_MORE will be set in flags on return to indicate that another t_rcvvudata() should be called to retrieve the rest of the data unit. Subsequent calls to t_rcvvudata() will return zero for the length of the address and options, until the full data unit has been received.</p>
<b>RETURN VALUES</b>	On successful completion, t_rcvvudata() returns the number of bytes received. Otherwise, it returns -1 on failure and t_errno is set to indicate the error.

**VALID STATES** T\_IDLE.

**ERRORS** On failure, `t_errno` is set to one of the following:

TBADDATA	<code>iovcount</code> is greater than <code>T_IOV_MAX</code> .
TBADF	The specified file descriptor does not refer to a transport endpoint.
TBUFOVFLW	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 <i>unitdata</i> will be discarded.
TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
TNODATA	<code>O_NONBLOCK</code> was set, but no data units are currently available from the transport provider.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI  
COMPATIBILITY  
ATTRIBUTES**

In the TLI interface definition, no counterpart of this routine was defined.

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `fcntl(2)`, `t_alloc(3NSL)`, `t_open(3NSL)`, `t_rcvvudata(3NSL)`, `t_rcvuderr(3NSL)`, `t_sndudata(3NSL)`, `t_sndvudata(3NSL)`, `attributes(5)`

## t\_snd(3NSL)

<b>NAME</b>	t_snd – send data or expedited data over a connection
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_snd(int fd, void *buf, unsigned int nbytes, int flags);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function is used to send either normal or expedited data. The argument <i>fd</i> identifies the local transport endpoint over which data should be sent, <i>buf</i> points to the user data, <i>nbytes</i> specifies the number of bytes of user data to be sent, and <i>flags</i> specifies any optional flags described below:</p> <p><b>T_EXPEDITED</b>     If set in <i>flags</i>, the data will be sent as expedited data and will be subject to the interpretations of the transport provider.</p> <p><b>T_MORE</b>            If set in <i>flags</i>, this indicates to the transport provider that the transport service data unit (TSDU) (or expedited transport service data unit - ETSDU) is being sent through multiple <code>t_snd()</code> calls. Each <code>t_snd()</code> with the <b>T_MORE</b> flag set indicates that another <code>t_snd()</code> will follow with more data for the current TSDU (or ETSDU).</p> <p>The end of the TSDU (or ETSDU) is identified by a <code>t_snd()</code> call with the <b>T_MORE</b> flag not set. Use of <b>T_MORE</b> 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 <i>info</i> argument on return from <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>, the <b>T_MORE</b> flag is not meaningful and will be ignored if set.</p> <p>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; that is, when the <b>T_MORE</b> flag is not set. Some transport providers also forbid zero-length TSDUs and ETSDUs.</p> <p><b>T_PUSH</b>            If set in <i>flags</i>, requests that the provider transmit all data that it has accumulated but not sent. The request is a local action on the provider and does not affect any similarly named protocol flag (for example, the TCP PUSH flag). This effect of setting this flag is protocol-dependent, and it may be ignored entirely by transport providers which do not support the use of this feature.</p> <p>Note that the communications provider is free to collect data in a send buffer until it accumulates a sufficient amount for transmission.</p>



By default, `t_snd()` 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 `O_NONBLOCK` is set by means of `t_open(3NSL)` or `fcntl(2)`, `t_snd()` will execute in asynchronous mode, and will fail immediately if there are flow control restrictions. The process can arrange to be informed when the flow control restrictions are cleared by means of either `t_look(3NSL)` or the EM interface.

On successful completion, `t_snd()` returns the number of bytes (octets) accepted by the communications provider. Normally this will equal the number of octets specified in `nbytes`. However, if `O_NONBLOCK` is set or the function is interrupted by a signal, it is possible that only part of the data has actually been accepted by the communications provider. In this case, `t_snd()` returns a value that is less than the value of `nbytes`. If `t_snd()` is interrupted by a signal before it could transfer data to the communications provider, it returns `-1` with `t_errno` set to `TSYSERR` and `errno` set to `EINTR`.

If `nbytes` is zero and sending of zero bytes is not supported by the underlying communications service, `t_snd()` returns `-1` with `t_errno` set to `TBADDATA`.

The size of each TSDU or ETSDU must not exceed the limits of the transport provider as specified by the current values in the TSDU or ETSDU fields in the *info* argument returned by `t_getinfo(3NSL)`.

The error `TLOOK` is returned for asynchronous events. It is required only for an incoming disconnect event but may be returned for other events.

## RETURN VALUES

On successful completion, `t_snd()` 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 if the number of bytes accepted by the communications provider is less than the number of bytes requested, this may either indicate that `O_NONBLOCK` is set and the communications provider is blocked due to flow control, or that `O_NONBLOCK` is clear and the function was interrupted by a signal.

## ERRORS

On failure, `t_errno` is set to one of the following:

- |                       |  |
|-----------------------|--|
| <code>TBADDATA</code> | 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.</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</li> </ul> |
|-----------------------|--|

t\_snd(3NSL)

	implementation to detect such an error case is implementation-dependent. See WARNINGS, below.
TBADF	The specified file descriptor does not refer to a transport endpoint.
TBADFLAG	An invalid flag was specified.
TFLOW	O_NONBLOCK 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 this transport endpoint.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TSYSERR	A system error has occurred during execution of this function.
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>	The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header:  <pre>#include &lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>	The <code>t_errno</code> values that can be set by the XTI interface and cannot be set by the TLI interface are:  TPROTO TLOOK TBADFLAG TOUTSTATE  The <code>t_errno</code> values that this routine can return under different circumstances than its XTI counterpart are:  TBADDATA  In the TBADDATA error cases described above, TBADDATA is returned, only for illegal zero byte TSDU (ETSDU) send attempts.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** fcntl(2), t\_getinfo(3NSL), t\_look(3NSL), t\_open(3NSL), t\_rcv(3NSL), attributes(5)

**WARNINGS** 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 XTI. 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 XTI, `t_snd()` fails with `TBADDATA`.

## t\_snddis(3NSL)

<b>NAME</b>	t_snddis – send user-initiated disconnection request								
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_snddis(int fd, const struct t_call *call);</pre>								
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function is used to initiate an abortive release on an already established connection, or to reject a connection request. The argument <i>fd</i> identifies the local transport endpoint of the connection, and <i>call</i> specifies information associated with the abortive release. The argument <i>call</i> points to a <code>t_call</code> structure which contains the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata; int sequence;</pre> <p>The values in <i>call</i> have different semantics, depending on the context of the call to <code>t_snddis()</code>. When rejecting a connection request, <i>call</i> must be non-null and contain a valid value of <i>sequence</i> to uniquely identify the rejected connection indication to the transport provider. The <i>sequence</i> field is only meaningful if the transport connection is in the <code>T_INCON</code> state. The <i>addr</i> and <i>opt</i> fields of <i>call</i> are ignored. In all other cases, <i>call</i> need only be used when data is being sent with the disconnection request. The <i>addr</i>, <i>opt</i> and <i>sequence</i> fields of the <code>t_call</code> structure are ignored. If the user does not wish to send data to the remote user, the value of <i>call</i> may be a null pointer.</p> <p>The <i>udata</i> structure 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 <i>discon</i> field, of the <i>info</i> argument of <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>. If the <i>len</i> field of <i>udata</i> is zero, no data will be sent to the remote user.</p>								
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.								
<b>VALID STATES</b>	<code>T_DATAXFER</code> , <code>T_OUTCON</code> , <code>T_OUTREL</code> , <code>T_INREL</code> , <code>T_INCON</code> ( <code>ocnt &gt; 0</code> ).								
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following:								
	<table><tr><td>TBADF</td><td>The specified file descriptor does not refer to a transport endpoint.</td></tr><tr><td>TBADDATA</td><td>The amount of user data specified was not within the bounds allowed by the transport provider.</td></tr><tr><td>TBADSEQ</td><td>An invalid sequence number was specified, or a null <i>call</i> pointer was specified, when rejecting a connection request.</td></tr><tr><td>TLOOK</td><td>An asynchronous event, which requires attention, has occurred.</td></tr></table>	TBADF	The specified file descriptor does not refer to a transport endpoint.	TBADDATA	The amount of user data specified was not within the bounds allowed by the transport provider.	TBADSEQ	An invalid sequence number was specified, or a null <i>call</i> pointer was specified, when rejecting a connection request.	TLOOK	An asynchronous event, which requires attention, has occurred.
TBADF	The specified file descriptor does not refer to a transport endpoint.								
TBADDATA	The amount of user data specified was not within the bounds allowed by the transport provider.								
TBADSEQ	An invalid sequence number was specified, or a null <i>call</i> pointer was specified, when rejecting a connection request.								
TLOOK	An asynchronous event, which requires attention, has occurred.								

TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI COMPATIBILITY**

The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.

**Interface Header**

The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Error Description Values**

The `t_errno` value TPROTO can be set by the XTI interface but not by the TLI interface.

**Option Buffers**

The format of the options in an `opt` buffer is dictated by the transport provider. Unlike the XTI interface, the TLI interface does not fix the buffer format.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO**

`t_connect(3NSL)`, `t_getinfo(3NSL)`, `t_listen(3NSL)`, `t_open(3NSL)`, `t_snd(3NSL)`, `attributes(5)`

**WARNINGS**

`t_snddis()` is an abortive disconnection. Therefore a `t_snddis()` issued on a connection endpoint may cause data previously sent by means of `t_snd(3NSL)`, or data not yet received, to be lost, even if an error is returned.

## t\_sndrel(3NSL)

<b>NAME</b>	t_sndrel – initiate an orderly release														
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_sndrel(int fd);</pre>														
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>For transport providers of type <code>T_COTS_ORD</code>, this function is used to initiate an orderly release of the outgoing direction of data transfer and indicates to the transport provider that the transport user has no more data to send. The argument <i>fd</i> identifies the local transport endpoint where the connection exists. After calling <code>t_sndrel()</code>, 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. For transport providers of types other than <code>T_COTS_ORD</code>, this function fails with error <code>TNOTSUPPORT</code>.</p>														
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.														
<b>VALID STATES</b>	<code>T_DATAXFER</code> , <code>T_INREL</code> .														
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following:														
	<table><tr><td><code>TBADF</code></td><td>The specified file descriptor does not refer to a transport endpoint.</td></tr><tr><td><code>TFLOW</code></td><td><code>O_NONBLOCK</code> was set, but the flow control mechanism prevented the transport provider from accepting the function at this time.</td></tr><tr><td><code>TLOOK</code></td><td>An asynchronous event has occurred on this transport endpoint and requires immediate attention.</td></tr><tr><td><code>TNOTSUPPORT</code></td><td>This function is not supported by the underlying transport provider.</td></tr><tr><td><code>TOUTSTATE</code></td><td>The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.</td></tr><tr><td><code>TPROTO</code></td><td>This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<code>t_errno</code>).</td></tr><tr><td><code>TSYSERR</code></td><td>A system error has occurred during execution of this function.</td></tr></table>	<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.	<code>TFLOW</code>	<code>O_NONBLOCK</code> was set, but the flow control mechanism prevented the transport provider from accepting the function at this time.	<code>TLOOK</code>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.	<code>TNOTSUPPORT</code>	This function is not supported by the underlying transport provider.	<code>TOUTSTATE</code>	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.	<code>TPROTO</code>	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).	<code>TSYSERR</code>	A system error has occurred during execution of this function.
<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.														
<code>TFLOW</code>	<code>O_NONBLOCK</code> was set, but the flow control mechanism prevented the transport provider from accepting the function at this time.														
<code>TLOOK</code>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.														
<code>TNOTSUPPORT</code>	This function is not supported by the underlying transport provider.														
<code>TOUTSTATE</code>	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.														
<code>TPROTO</code>	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).														
<code>TSYSERR</code>	A system error has occurred during execution of this function.														
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.														

t\_sndrel(3NSL)

**Interface Header** The XTI interfaces use the header file, `xti.h`. TLI interfaces should *not* use this header. They should use the header:

```
#include <tiuser.h>
```

**Error Description Values** The `t_errno` values that can be set by the XTI interface and cannot be set by the TLI interface are:

```
TPROTO  
TLOOK  
TOUTSTATE
```

**Notes** Whenever this function fails with `t_error` set to `TFLOW`, `O_NONBLOCK` must have been set.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `t_error(3NSL)`, `t_getinfo(3NSL)`, `t_open(3NSL)`, `t_rcvrel(3NSL)`, `attributes(5)`

## t\_sndreldata(3NSL)

<b>NAME</b>	t_sndreldata – initiate or respond to an orderly release with user data				
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_sndreldata(int fd, struct t_discon *discon);</pre>				
<b>DESCRIPTION</b>	<p>This function is used to initiate an orderly release of the outgoing direction of data transfer and to send user data with the release. The argument <i>fd</i> identifies the local transport endpoint where the connection exists, and <i>discon</i> points to a <code>t_discon</code> structure containing the following members:</p> <pre>struct netbuf udata; int reason; int sequence;</pre> <p>After calling <code>t_sndreldata()</code>, 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.</p> <p>The field <i>reason</i> specifies the reason for the disconnection through a protocol-dependent <i>reason code</i>, and <i>udata</i> identifies any user data that is sent with the disconnection; the field <i>sequence</i> is not used.</p> <p>The <i>udata</i> structure 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 <i>discon</i> field of the <i>info</i> argument of <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>. If the <i>len</i> field of <i>udata</i> is zero or if the provider did not return <code>T_ORDRELDATA</code> in the <code>t_open(3NSL)</code> flags, no data will be sent to the remote user.</p> <p>If a user does not wish to send data and reason code to the remote user, the value of <i>discon</i> may be a null pointer.</p> <p>This function is an optional service of the transport provider, only supported by providers of service type <code>T_COTS_ORD</code>. The flag <code>T_ORDRELDATA</code> in the <i>info</i>→<i>flag</i> field returned by <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code> indicates that the provider supports orderly release user data.</p> <p>This function may not be available on all systems.</p>				
<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.				
<b>VALID STATES</b>	<code>T_DATAXFER</code> , <code>T_INREL</code> .				
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following:				
	<table><tr><td><code>TBADDATA</code></td><td>The amount of user data specified was not within the bounds allowed by the transport provider, or user data was supplied and the provider did not return <code>T_ORDRELDATA</code> in the <code>t_open(3NSL)</code> flags.</td></tr><tr><td><code>TBADF</code></td><td>The specified file descriptor does not refer to a transport endpoint.</td></tr></table>	<code>TBADDATA</code>	The amount of user data specified was not within the bounds allowed by the transport provider, or user data was supplied and the provider did not return <code>T_ORDRELDATA</code> in the <code>t_open(3NSL)</code> flags.	<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.
<code>TBADDATA</code>	The amount of user data specified was not within the bounds allowed by the transport provider, or user data was supplied and the provider did not return <code>T_ORDRELDATA</code> in the <code>t_open(3NSL)</code> flags.				
<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.				



t\_sndreldata(3NSL)

TFLOW	O_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting the function at this time.
TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
TNOTSUPPORT	Orderly release is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <i>t_errno</i> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI  
COMPATIBILITY  
ATTRIBUTES**

In the TLI interface definition, no counterpart of this routine was defined.

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `t_getinfo(3NSL)`, `t_open(3NSL)`, `t_rcvrel(3NSL)`, `t_rcvreldata(3NSL)`, `t_sndrel(3NSL)`, `attributes(5)`

**NOTES** The interfaces `t_sndreldata()` and `t_rcvreldata(3NSL)` are only for use with a specific transport called "minimal OSI," which is not available on the Solaris platform. These interfaces are not available for use in conjunction with Internet Transports (TCP or UDP).

t\_sndudata(3NSL)

<b>NAME</b>	t_sndudata – send a data unit
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_sndudata(int fd, const struct t_unitdata *unitdata);</pre>
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>This function is used in connectionless-mode to send a data unit to another transport user. The argument <i>fd</i> identifies the local transport endpoint through which data will be sent, and <i>unitdata</i> points to a <code>t_unitdata</code> structure containing the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata;</pre> <p>In <i>unitdata</i>, <i>addr</i> specifies the protocol address of the destination user, <i>opt</i> identifies options that the user wants associated with this request, and <i>udata</i> specifies the user data to be sent. The user may choose not to specify what protocol options are associated with the transfer by setting the <i>len</i> field of <i>opt</i> to zero. In this case, the provider uses the option values currently set for the communications endpoint.</p> <p>If the <i>len</i> field of <i>udata</i> is zero, and sending of zero octets is not supported by the underlying transport service, the <code>t_sndudata()</code> will return <code>-1</code> with <code>t_errno</code> set to <code>TBADDATA</code>.</p> <p>By default, <code>t_sndudata()</code> 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 <code>O_NONBLOCK</code> is set by means of <code>t_open(3NSL)</code> or <code>fcntl(2)</code>, <code>t_sndudata()</code> will execute in asynchronous mode and will fail under such conditions. The process can arrange to be notified of the clearance of a flow control restriction by means of either <code>t_look(3NSL)</code> or the EM interface.</p> <p>If the amount of data specified in <i>udata</i> exceeds the TSDU size as returned in the <i>tsdu</i> field of the <i>info</i> argument of <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>, a <code>TBADDATA</code> error will be generated. If <code>t_sndudata()</code> is called before the destination user has activated its transport endpoint (see <code>t_bind(3NSL)</code>), the data unit may be discarded.</p> <p>If it is not possible for the transport provider to immediately detect the conditions that cause the errors <code>TBADDADDR</code> and <code>TBADOPT</code>, these errors will alternatively be returned by <code>t_rcvuderr</code>. Therefore, an application must be prepared to receive these errors in both of these ways.</p> <p>If the call is interrupted, <code>t_sndudata()</code> will return <code>EINTR</code> and the datagram will not be sent.</p>

<b>RETURN VALUES</b>	Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and <code>t_errno</code> is set to indicate an error.
<b>VALID STATES</b>	<code>T_IDLE</code> .
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following: <ul style="list-style-type: none"> <li><code>TBADADDR</code>      The specified protocol address was in an incorrect format or contained illegal information.</li> <li><code>TBADDATA</code>      Illegal amount of data. A single send was attempted specifying a TSDU greater than that specified in the <i>info</i> argument, or a send of a zero byte TSDU is not supported by the provider.</li> <li><code>TBADF</code>            The specified file descriptor does not refer to a transport endpoint.</li> <li><code>TBADOPT</code>        The specified options were in an incorrect format or contained illegal information.</li> <li><code>TFLOW</code>            <code>O_NONBLOCK</code> was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.</li> <li><code>TLOOK</code>            An asynchronous event has occurred on this transport endpoint.</li> <li><code>TNOTSUPPORT</code>    This function is not supported by the underlying transport provider.</li> <li><code>TOUTSTATE</code>      The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.</li> <li><code>TPROTO</code>           This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<code>t_errno</code>).</li> <li><code>TSYSERR</code>        A system error has occurred during execution of this function.</li> </ul>
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.
<b>Interface Header</b>	The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header: <pre>#include &lt;tiuser.h&gt;</pre>
<b>Error Description Values</b>	The <code>t_errno</code> values that can be set by the XTI interface and cannot be set by the TLI interface are: <ul style="list-style-type: none"> <li><code>TPROTO</code></li> <li><code>TBADADDR</code></li> <li><code>TBADOPT</code></li> <li><code>TLOOK</code></li> </ul>

## t\_sndudata(3NSL)

TOUTSTATE

**Notes** Whenever this function fails with `t_error` set to `TFLOW`, `O_NONBLOCK` must have been set.

**Option Buffers** The format of the options in an `opt` buffer is dictated by the transport provider. Unlike the XTI interface, the TLI interface does not fix the buffer format.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

**SEE ALSO** `fcntl(2)`, `t_alloc(3NSL)`, `t_bind(3NSL)`, `t_error(3NSL)`, `t_getinfo(3NSL)`, `t_look(3NSL)`, `t_open(3NSL)`, `t_rcvudata(3NSL)`, `t_rcvuderr(3NSL)`, `attributes(5)`

<b>NAME</b>	t_sndv – send data or expedited data, from one or more non-contiguous buffers, on a connection
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_sndv(int fd, const struct t_iovec *iov, unsigned int iovcount,            int flags);</pre>
<b>DESCRIPTION</b>	<p>This function is used to send either normal or expedited data. The argument <i>fd</i> identifies the local transport endpoint over which data should be sent, <i>iov</i> points to an array of buffer address/buffer length pairs. <code>t_sndv()</code> sends data contained in buffers <i>iov0</i>, <i>iov1</i>, through <i>iov</i> [<i>iovcount-1</i>]. <i>iovcount</i> contains the number of non-contiguous data buffers which is limited to <code>T_IOV_MAX</code>, an implementation-defined value of at least 16. If the limit is exceeded, the function fails with <code>TBADDATA</code>.</p> <p><i>iov(0).iov_len + .. + iov(iovcount-1).iov_len</i></p> <p>Note that the limit on the total number of bytes available in all buffers passed: may be constrained by implementation limits. If no other constraint applies, it will be limited by <code>INT_MAX</code>. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.</p> <p>The argument <i>flags</i> specifies any optional flags described below:</p> <p><code>T_EXPEDITED</code>      If set in <i>flags</i>, the data will be sent as expedited data and will be subject to the interpretations of the transport provider.</p> <p><code>T_MORE</code>            If set in <i>flags</i>, this indicates to the transport provider that the transport service data unit (TSDU) (or expedited transport service data unit – ETSDU) is being sent through multiple <code>t_sndv()</code> calls. Each <code>t_sndv()</code> with the <code>T_MORE</code> flag set indicates that another <code>t_sndv()</code> or <code>t_snd(3NSL)</code> will follow with more data for the current TSDU (or ETSDU).</p> <p>The end of the TSDU (or ETSDU) is identified by a <code>t_sndv()</code> call with the <code>T_MORE</code> flag not set. Use of <code>T_MORE</code> 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 <i>info</i> argument on return from <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>, the <code>T_MORE</code> flag is not meaningful and will be ignored if set.</p> <p>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, that is, when the <code>T_MORE</code> flag is not set. Some transport providers also forbid zero-length TSDUs and ETSDUs.</p>

## t\_sndv(3NSL)

If set in *flags*, requests that the provider transmit all data that it has accumulated but not sent. The request is a local action on the provider and does not affect any similarly named protocol flag (for example, the TCP PUSH flag). This effect of setting this flag is protocol-dependent, and it may be ignored entirely by transport providers which do not support the use of this feature.

The communications provider is free to collect data in a send buffer until it accumulates a sufficient amount for transmission.

By default, `t_sndv()` 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 `O_NONBLOCK` is set by means of `t_open(3NSL)` or `fcntl(2)`, `t_sndv()` executes in asynchronous mode, and will fail immediately if there are flow control restrictions. The process can arrange to be informed when the flow control restrictions are cleared via either `t_look(3NSL)` or the EM interface.

On successful completion, `t_sndv()` returns the number of bytes accepted by the transport provider. Normally this will equal the total number of bytes to be sent, that is,

```
(iov0.iov_len + .. + iov[iovcnt-1].iov_len)
```

However, the interface is constrained to send at most `INT_MAX` bytes in a single send. When `t_sndv()` has submitted `INT_MAX` (or lower constrained value, see the note above) bytes to the provider for a single call, this value is returned to the user. However, if `O_NONBLOCK` is set or the function is interrupted by a signal, it is possible that only part of the data has actually been accepted by the communications provider. In this case, `t_sndv()` returns a value that is less than the value of `nbytes`. If `t_sndv()` is interrupted by a signal before it could transfer data to the communications provider, it returns `-1` with `t_errno` set to `TSYSERR` and `errno` set to `EINTR`.

If the number of bytes of data in the *iov* array is zero and sending of zero octets is not supported by the underlying transport service, `t_sndv()` returns `-1` with `t_errno` set to `TBADDATA`.

The size of each TSDU or ETSDU must not exceed the limits of the transport provider as specified by the current values in the TSDU or ETSDU fields in the *info* argument returned by `t_getinfo(3NSL)`.

The error `TLOOK` is returned for asynchronous events. It is required only for an incoming disconnect event but may be returned for other events.

### RETURN VALUES

On successful completion, `t_sndv()` 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 synchronous mode, if more than `INT_MAX` bytes of data are passed in the *iov* array, only the first `INT_MAX` bytes will be passed to the provider.

If the number of bytes accepted by the communications provider is less than the number of bytes requested, this may either indicate that `O_NONBLOCK` is set and the communications provider is blocked due to flow control, or that `O_NONBLOCK` is clear and the function was interrupted by a signal.

**VALID STATES**

T\_DATAXFER, T\_INREL.

**ERRORS**

On failure, `t_errno` is set to one of the following:

TBADDATA	Illegal amount of data:
TBADF	The specified file descriptor does not refer to a transport endpoint. <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.</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 <code>WARNINGS</code>, below.</li> <li>■ <i>iovcount</i> is greater than <code>T_IOV_MAX</code>.</li> </ul>
TBADFLAG	An invalid flag was specified.
TFLOW	<code>O_NONBLOCK</code> 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 this transport endpoint.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).
TSYSERR	A system error has occurred during execution of this function.

**TLI  
COMPATIBILITY  
ATTRIBUTES**

In the TLI interface definition, no counterpart of this routine was defined.

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

t\_sndv(3NSL)

**SEE ALSO** t\_getinfo(3NSL), t\_open(3NSL), t\_rcvv(3NSL) t\_rcv(3NSL), t\_snd(3NSL), attributes(5)

**WARNINGS** 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\_sndv() or t\_snd(3NSL) calls, then the different data may be intermixed.

Multiple sends which exceed the maximum TSDU or ETSDU size may not be discovered by XTI. 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 TSYSEERR, a TBADDDATA or a TPROTO error.

If multiple sends which exceed the maximum TSDU or ETSDU size are detected by XTI, t\_sndv() fails with TBADDDATA.



<b>NAME</b>	t_sndvudata – send a data unit from one or more noncontiguous buffers
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_sndvudata(int fd, struct t_unitdata *unitdata, struct t_iovec                 *iov, unsigned int iovcount);</pre>
<b>DESCRIPTION</b>	<p>This function is used in connectionless mode to send a data unit to another transport user. The argument <i>fd</i> identifies the local transport endpoint through which data will be sent, <i>iovcount</i> contains the number of non-contiguous <i>udata</i> buffers and is limited to an implementation-defined value given by <code>T_IOV_MAX</code> which is at least 16, and <i>unitdata</i> points to a <code>t_unitdata</code> structure containing the following members:</p> <pre>struct netbuf addr; struct netbuf opt; struct netbuf udata;</pre> <p>If the limit on <i>iovcount</i> is exceeded, the function fails with <code>TBADDATA</code>.</p> <p>In <i>unitdata</i>, <i>addr</i> specifies the protocol address of the destination user, and <i>opt</i> identifies options that the user wants associated with this request. The <i>udata</i> field is not used. The user may choose not to specify what protocol options are associated with the transfer by setting the <i>len</i> field of <i>opt</i> to zero. In this case, the provider may use default options.</p> <p>The data to be sent is identified by <i>iov</i>[0] through <i>iov</i>[<i>iovcount</i>-1].</p> <p>Note that the limit on the total number of bytes available in all buffers passed:</p> <pre>iov(0).iov_len + . . . + iov(iovcount-1).iov_len</pre> <p>may be constrained by implementation limits. If no other constraint applies, it will be limited by <code>INT_MAX</code>. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.</p> <p>By default, <code>t_sndvudata()</code> 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 <code>O_NONBLOCK</code> is set by means of <code>t_open(3NSL)</code> or <code>fcntl(2)</code>, <code>t_sndvudata()</code> executes in asynchronous mode and will fail under such conditions. The process can arrange to be notified of the clearance of a flow control restriction by means of either <code>t_look(3NSL)</code> or the EM interface.</p> <p>If the amount of data specified in <i>iov</i>[0] through <i>iov</i>[<i>iovcount</i>-1] exceeds the TSDU size as returned in the <i>tsdu</i> field of the <i>info</i> argument of <code>t_open(3NSL)</code> or <code>t_getinfo(3NSL)</code>, or is zero and sending of zero octets is not supported by the underlying transport service, a <code>TBADDATA</code> error is generated. If <code>t_sndvudata()</code> is called before the destination user has activated its transport endpoint (see <code>t_bind(3NSL)</code>), the data unit may be discarded.</p>

t\_sndvudata(3NSL)

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 t\_rcvuderr(3NSL). An application must therefore be prepared to receive these errors in both of these ways.

**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.

**VALID STATES** T\_IDLE.

**ERRORS** On failure, t\_errno is set to one of the following:

TBADADDR	The specified protocol address was in an incorrect format or contained illegal information.
TBADDATA	Illegal amount of data. <ul style="list-style-type: none"><li>■ A single send was attempted specifying a TSDU greater than that specified in the <i>info</i> argument, or a send of a zero byte TSDU is not supported by the provider.</li><li>■ <i>iovcount</i> is greater than T_IOV_MAX.</li></ul>
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	O_NONBLOCK i 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 this transport endpoint.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (t_errno).
TSYSERR	A system error has occurred during execution of this function.

**TLI  
COMPATIBILITY  
ATTRIBUTES**

In the TLI interface definition, no counterpart of this routine was defined.

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

t\_sndvudata(3NSL)

**SEE ALSO** fcntl(2), t\_alloc(3NSL), t\_open(3NSL), t\_rcvudata(3NSL),  
t\_rcvvudata(3NSL), t\_rcvuderr(3NSL), t\_sndudata(3NSL), attributes(5)

## t\_strerror(3NSL)

<b>NAME</b>	t_strerror – produce an error message string				
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  const char *t_strerror(int errnum);</pre>				
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>The <code>t_strerror()</code> function maps the error number in <code>errnum</code> that corresponds to an XTI error to a language-dependent error message string and returns a pointer to the string. The string pointed to will not be modified by the program, but may be overwritten by a subsequent call to the <code>t_strerror</code> function. The string is not terminated by a newline character. The language for error message strings written by <code>t_strerror()</code> is that of the current locale. If it is English, the error message string describing the value in <code>t_errno</code> may be derived from the comments following the <code>t_errno</code> codes defined in <code>&lt;xti.h&gt;</code>. If an error code is unknown, and the language is English, <code>t_strerror()</code> returns the string:</p> <pre>"&lt;error&gt;: error unknown" where &lt;error&gt; is the error number supplied as input. In other languages, an equivalent text is provided.</pre>				
<b>VALID STATES</b>	ALL - apart from T_UNINIT.				
<b>RETURN VALUES</b>	The function <code>t_strerror()</code> returns a pointer to the generated message string.				
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.				
<b>Interface Header</b>	The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header:				
	<pre>#include &lt;tiuser.h&gt;</pre>				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT Level	Safe				
<b>SEE ALSO</b>	<code>t_errno(3NSL)</code> , <code>t_error(3NSL)</code> , <code>attributes(5)</code>				

<b>NAME</b>	t_sync – synchronize transport library												
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_sync (int fd) ;</pre>												
<b>DESCRIPTION</b>	<p>This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>For the transport endpoint specified by <i>fd</i>, <code>t_sync()</code> synchronizes the data structures managed by the transport library with information from the underlying transport provider. In doing so, it can convert an uninitialized file descriptor (obtained by means of a <code>open(2)</code>, <code>dup(2)</code> or as a result of a <code>fork(2)</code> and <code>exec(2)</code>) to an initialized transport endpoint, assuming that the file descriptor referenced a transport endpoint, by updating and allocating the necessary library data structures. This function also allows two cooperating processes to synchronize their interaction with a transport provider.</p> <p>For example, if a process forks a new process and issues an <code>exec(2)</code>, the new process must issue a <code>t_sync()</code> to build the private library data structure associated with a transport endpoint and to synchronize the data structure with the relevant provider information.</p> <p>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 function <code>t_sync()</code> 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 <i>after</i> a <code>t_sync()</code> is issued.</p> <p>If the transport endpoint is undergoing a state transition when <code>t_sync()</code> is called, the function will fail.</p>												
<b>RETURN VALUES</b>	<p>On successful completion, the state of the transport endpoint is returned. Otherwise, a value of <code>-1</code> is returned and <code>t_errno</code> is set to indicate an error. The state returned is one of the following:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>T_UNBND</code></td> <td>Unbound.</td> </tr> <tr> <td><code>T_IDLE</code></td> <td>Idle.</td> </tr> <tr> <td><code>T_OUTCON</code></td> <td>Outgoing connection pending.</td> </tr> <tr> <td><code>T_INCON</code></td> <td>Incoming connection pending.</td> </tr> <tr> <td><code>T_DATAXFER</code></td> <td>Data transfer.</td> </tr> <tr> <td><code>T_OUTREL</code></td> <td>Outgoing orderly release (waiting for an orderly release indication).</td> </tr> </table>	<code>T_UNBND</code>	Unbound.	<code>T_IDLE</code>	Idle.	<code>T_OUTCON</code>	Outgoing connection pending.	<code>T_INCON</code>	Incoming connection pending.	<code>T_DATAXFER</code>	Data transfer.	<code>T_OUTREL</code>	Outgoing orderly release (waiting for an orderly release indication).
<code>T_UNBND</code>	Unbound.												
<code>T_IDLE</code>	Idle.												
<code>T_OUTCON</code>	Outgoing connection pending.												
<code>T_INCON</code>	Incoming connection pending.												
<code>T_DATAXFER</code>	Data transfer.												
<code>T_OUTREL</code>	Outgoing orderly release (waiting for an orderly release indication).												

## t\_sync(3NSL)

	T_INREL	Incoming orderly release (waiting for an orderly release request).				
<b>ERRORS</b>	On failure, <code>t_errno</code> is set to one of the following:					
	TBADF	The specified file descriptor does not refer to a transport endpoint. This error may be returned when the <i>fd</i> has been previously closed or an erroneous number may have been passed to the call.				
	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <code>t_errno</code> ).				
	TSTATECHNG	The transport endpoint is undergoing a state change.				
	TSYSERR	A system error has occurred during execution of this function.				
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.					
<b>Interface Header</b>	The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header:					
	<pre>#include &lt;tiuser.h&gt;</pre>					
<b>Error Description Values</b>	The <code>t_errno</code> value that can be set by the XTI interface and cannot be set by the TLI interface is:					
	TPROTO					
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:					
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT Level</td><td>Safe</td></tr></tbody></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
MT Level	Safe					
<b>SEE ALSO</b>	<code>dup(2)</code> , <code>exec(2)</code> , <code>fork(2)</code> , <code>open(2)</code> , <code>attributes(5)</code>					

<b>NAME</b>	t_sysconf – get configurable XTI variables				
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_sysconf(intname);</pre>				
<b>DESCRIPTION</b>	<p>The t_sysconf() function provides a method for the application to determine the current value of configurable and implementation-dependent XTI limits or options.</p> <p>The <i>name</i> argument represents the XTI system variable to be queried. The following table lists the minimal set of XTI system variables from &lt;xti.h&gt; that can be returned by t_sysconf(), and the symbolic constants, defined in &lt;xti.h&gt; that are the corresponding values used for <i>name</i>.</p> <table border="1" data-bbox="461 720 1430 810"> <thead> <tr> <th>Variable</th> <th>Value of Name</th> </tr> </thead> <tbody> <tr> <td>T_IOV_MAX</td> <td>_SC_T_IOV_MAX</td> </tr> </tbody> </table>	Variable	Value of Name	T_IOV_MAX	_SC_T_IOV_MAX
Variable	Value of Name				
T_IOV_MAX	_SC_T_IOV_MAX				
<b>RETURN VALUES</b>	If <i>name</i> is valid, t_sysconf() returns the value of the requested limit/option, which might be -1, and leaves t_errno unchanged. Otherwise, a value of -1 is returned and t_errno is set to indicate an error.				
<b>VALID STATES</b>	All.				
<b>ERRORS</b>	On failure, t_errno is set to the following: TBADFLAG <i>name</i> has an invalid value.				
<b>TLI COMPATIBILITY ATTRIBUTES</b>	In the TLI interface definition, no counterpart of this routine was defined. See attributes(5) for descriptions of the following attributes:				
	<table border="1" data-bbox="461 1262 1430 1352"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
<b>SEE ALSO</b>	sysconf(3C), t_rcvv(3NSL), t_rcvvdata(3NSL), t_sndv(3NSL), t_sndvdata(3NSL), attributes(5)				

## t\_unbind(3NSL)

<b>NAME</b>	t_unbind – disable a transport endpoint										
<b>SYNOPSIS</b>	<pre>#include &lt;xti.h&gt;  int t_unbind(int fd);</pre>										
<b>DESCRIPTION</b>	<p>The This routine is part of the XTI interfaces which evolved from the TLI interfaces. XTI represents the future evolution of these interfaces. However, TLI interfaces are supported for compatibility. When using a TLI routine that has the same name as an XTI routine, the <code>tiuser.h</code> header file must be used. Refer to the TLI COMPATIBILITY section for a description of differences between the two interfaces.</p> <p>t_unbind() function disables the transport endpoint specified by <i>fd</i> which was previously bound by t_bind(3NSL). On completion of this call, no further data or events destined for this transport endpoint will be accepted by the transport provider. An endpoint which is disabled by using t_unbind() can be enabled by a subsequent call to t_bind(3NSL).</p>										
<b>RETURN VALUES</b>	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.										
<b>VALID STATES</b>	T_IDLE.										
<b>ERRORS</b>	On failure, t_errno is set to one of the following:										
	<table><tr><td>TBADF</td><td>The specified file descriptor does not refer to a transport endpoint.</td></tr><tr><td>TLOOK</td><td>An asynchronous event has occurred on this transport endpoint.</td></tr><tr><td>TOUTSTATE</td><td>The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.</td></tr><tr><td>TPROTO</td><td>This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (t_errno).</td></tr><tr><td>TSYSERR</td><td>A system error has occurred during execution of this function.</td></tr></table>	TBADF	The specified file descriptor does not refer to a transport endpoint.	TLOOK	An asynchronous event has occurred on this transport endpoint.	TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.	TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (t_errno).	TSYSERR	A system error has occurred during execution of this function.
TBADF	The specified file descriptor does not refer to a transport endpoint.										
TLOOK	An asynchronous event has occurred on this transport endpoint.										
TOUTSTATE	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.										
TPROTO	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (t_errno).										
TSYSERR	A system error has occurred during execution of this function.										
<b>TLI COMPATIBILITY</b>	The XTI and TLI interface definitions have common names but use different header files. This, and other semantic differences between the two interfaces are described in the subsections below.										
<b>Interface Header</b>	The XTI interfaces use the header file, <code>xti.h</code> . TLI interfaces should <i>not</i> use this header. They should use the header:										
	<pre>#include &lt;tiuser.h&gt;</pre>										
<b>Error Description Values</b>	The t_errno value that can be set by the XTI interface and cannot be set by the TLI interface is:										
	TPROTO										



t\_unbind(3NSL)

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** t\_bind(3NSL), attributes(5)

## xdr(3NSL)

<b>NAME</b>	xdr – library routines for external data representation																																																				
<b>DESCRIPTION</b>	XDR routines allow C programmers to describe arbitrary data structures in a machine-independent fashion. Data for remote procedure calls (RPC) are transmitted using these routines.																																																				
<b>Index to Routines</b>	The following table lists XDR routines and the manual reference pages on which they are described:																																																				
	<table> <thead> <tr> <th>XDR Routine</th> <th>Manual Reference Page</th> </tr> </thead> <tbody> <tr> <td>xdr_array</td> <td>xdr_complex(3NSL)</td> </tr> <tr> <td>xdr_bool</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_bytes</td> <td>xdr_complex(3NSL)</td> </tr> <tr> <td>xdr_char</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_control</td> <td>xdr_admin(3NSL)</td> </tr> <tr> <td>xdr_destroy</td> <td>xdr_create(3NSL)</td> </tr> <tr> <td>xdr_double</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_enum</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_float</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_free</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_getpos</td> <td>xdr_admin(3NSL)</td> </tr> <tr> <td>xdr_hyper</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_inline</td> <td>xdr_admin(3NSL)</td> </tr> <tr> <td>xdr_int</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_long</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_longlong_t</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_opaque</td> <td>xdr_complex(3NSL)</td> </tr> <tr> <td>xdr_pointer</td> <td>xdr_complex(3NSL)</td> </tr> <tr> <td>xdr_quadruple</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_reference</td> <td>xdr_complex(3NSL)</td> </tr> <tr> <td>xdr_setpos</td> <td>xdr_admin(3NSL)</td> </tr> <tr> <td>xdr_short</td> <td>xdr_simple(3NSL)</td> </tr> <tr> <td>xdr_sizeof</td> <td>xdr_admin(3NSL)</td> </tr> <tr> <td>xdr_string</td> <td>xdr_complex(3NSL)</td> </tr> <tr> <td>xdr_u_char</td> <td>xdr_simple(3NSL)</td> </tr> </tbody> </table>	XDR Routine	Manual Reference Page	xdr_array	xdr_complex(3NSL)	xdr_bool	xdr_simple(3NSL)	xdr_bytes	xdr_complex(3NSL)	xdr_char	xdr_simple(3NSL)	xdr_control	xdr_admin(3NSL)	xdr_destroy	xdr_create(3NSL)	xdr_double	xdr_simple(3NSL)	xdr_enum	xdr_simple(3NSL)	xdr_float	xdr_simple(3NSL)	xdr_free	xdr_simple(3NSL)	xdr_getpos	xdr_admin(3NSL)	xdr_hyper	xdr_simple(3NSL)	xdr_inline	xdr_admin(3NSL)	xdr_int	xdr_simple(3NSL)	xdr_long	xdr_simple(3NSL)	xdr_longlong_t	xdr_simple(3NSL)	xdr_opaque	xdr_complex(3NSL)	xdr_pointer	xdr_complex(3NSL)	xdr_quadruple	xdr_simple(3NSL)	xdr_reference	xdr_complex(3NSL)	xdr_setpos	xdr_admin(3NSL)	xdr_short	xdr_simple(3NSL)	xdr_sizeof	xdr_admin(3NSL)	xdr_string	xdr_complex(3NSL)	xdr_u_char	xdr_simple(3NSL)
XDR Routine	Manual Reference Page																																																				
xdr_array	xdr_complex(3NSL)																																																				
xdr_bool	xdr_simple(3NSL)																																																				
xdr_bytes	xdr_complex(3NSL)																																																				
xdr_char	xdr_simple(3NSL)																																																				
xdr_control	xdr_admin(3NSL)																																																				
xdr_destroy	xdr_create(3NSL)																																																				
xdr_double	xdr_simple(3NSL)																																																				
xdr_enum	xdr_simple(3NSL)																																																				
xdr_float	xdr_simple(3NSL)																																																				
xdr_free	xdr_simple(3NSL)																																																				
xdr_getpos	xdr_admin(3NSL)																																																				
xdr_hyper	xdr_simple(3NSL)																																																				
xdr_inline	xdr_admin(3NSL)																																																				
xdr_int	xdr_simple(3NSL)																																																				
xdr_long	xdr_simple(3NSL)																																																				
xdr_longlong_t	xdr_simple(3NSL)																																																				
xdr_opaque	xdr_complex(3NSL)																																																				
xdr_pointer	xdr_complex(3NSL)																																																				
xdr_quadruple	xdr_simple(3NSL)																																																				
xdr_reference	xdr_complex(3NSL)																																																				
xdr_setpos	xdr_admin(3NSL)																																																				
xdr_short	xdr_simple(3NSL)																																																				
xdr_sizeof	xdr_admin(3NSL)																																																				
xdr_string	xdr_complex(3NSL)																																																				
xdr_u_char	xdr_simple(3NSL)																																																				

xdr_u_hyper	xdr_simple(3NSL)
xdr_u_int	xdr_simple(3NSL)
xdr_u_long	xdr_simple(3NSL)
xdr_u_longlong_t	xdr_simple(3NSL)
xdr_u_short	xdr_simple(3NSL)
xdr_union	xdr_complex(3NSL)
xdr_vector	xdr_complex(3NSL)
xdr_void	xdr_simple(3NSL)
xdr_wrapstring	xdr_complex(3NSL)
xdrmem_create	xdr_create(3NSL)
xdrrec_create	xdr_create(3NSL)
xdrrec_endofrecord	xdr_admin(3NSL)
xdrrec_eof	xdr_admin(3NSL)
xdrrec_readbytes	xdr_admin(3NSL)
xdrrec_skiprecord	xdr_admin(3NSL)
xdrstdio_create	xdr_create(3NSL)

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** rpc(3NSL), xdr\_admin(3NSL), xdr\_complex(3NSL), xdr\_create(3NSL), xdr\_simple(3NSL), attributes(5)

## xdr\_admin(3NSL)

<b>NAME</b>	xdr_admin, xdr_control, xdr_getpos, xdr_inline, xdrrec_endofrecord, xdrrec_eof, xdrrec_readbytes, xdrrec_skiprecord, xdr_setpos, xdr_sizeof – library routines for external data representation
<b>DESCRIPTION</b>	<p>XDR library routines allow C programmers to describe arbitrary data structures in a machine-independent fashion. Protocols such as remote procedure calls (RPC) use these routines to describe the format of the data.</p> <p>These routines deal specifically with the management of the XDR stream.</p>
<b>Routines</b>	<p>See <code>rpc(3NSL)</code> for the definition of the XDR data structure. Note that any buffers passed to the XDR routines must be properly aligned. It is suggested either that <code>malloc(3C)</code> be used to allocate these buffers, or that the programmer insure that the buffer address is divisible evenly by four.</p> <pre>#include &lt;rpc/xdr.h&gt;</pre> <p><code>bool_t xdr_control( XDR *xdrs, int req, void *info );</code> A function macro to change or retrieve various information about an XDR stream. <i>req</i> indicates the type of operation and <i>info</i> is a pointer to the information. The supported values of <i>req</i> is <code>XDR_GET_BYTES_AVAIL</code> and its argument type is <code>xdr_bytesrec *</code>. They return the number of bytes left unconsumed in the stream and a flag indicating whether or not this is the last fragment.</p> <p><code>uint_t xdr_getpos( const XDR *xdrs );</code> A macro that invokes the get-position routine associated with the XDR stream, <i>xdrs</i>. The routine returns an unsigned integer, which indicates the position of the XDR byte stream. A desirable feature of XDR streams is that simple arithmetic works with this number, although the XDR stream instances need not guarantee this. Therefore, applications written for portability should not depend on this feature.</p> <p><code>long *xdr_inline( XDR *xdrs, const int len );</code> A macro that invokes the in-line routine associated with the XDR stream, <i>xdrs</i>. The routine returns a pointer to a contiguous piece of the stream's buffer; <i>len</i> is the byte length of the desired buffer. Note: pointer is cast to <code>long *</code>.</p> <p>Warning: <code>xdr_inline()</code> may return <code>NULL (0)</code> if it cannot allocate a contiguous piece of a buffer. Therefore the behavior may vary among stream instances; it exists for the sake of efficiency, and applications written for portability should not depend on this feature.</p> <p><code>bool_t xdrrec_endofrecord( XDR *xdrs, int sendnow );</code> This routine can be invoked only on streams created by <code>xdrrec_create()</code>. See <code>xdr_create(3NSL)</code>. The data in the output buffer is marked as a completed record, and the output buffer is optionally written out if <i>sendnow</i> is non-zero. This routine returns <code>TRUE</code> if it succeeds, <code>FALSE</code> otherwise.</p> <p><code>bool_t xdrrec_eof( XDR *xdrs );</code> This routine can be invoked only on streams created by <code>xdrrec_create()</code>. After consuming the rest of the current record in the stream, this routine returns <code>TRUE</code> if</p>

there is no more data in the stream's input buffer. It returns `FALSE` if there is additional data in the stream's input buffer.

```
int xdrrec_readbytes(XDR *xdrs, caddr_t addr, uint_t nbytes);
```

This routine can be invoked only on streams created by `xdrrec_create()`. It attempts to read *nbytes* bytes from the XDR stream into the buffer pointed to by *addr*. Upon success this routine returns the number of bytes read. Upon failure, it returns `-1`. A return value of `0` indicates an end of record.

```
bool_t xdrrec_skiprecord(XDR *xdrs);
```

This routine can be invoked only on streams created by `xdrrec_create()`. See `xdr_create(3NSL)`. It tells the XDR implementation that the rest of the current record in the stream's input buffer should be discarded. This routine returns `TRUE` if it succeeds, `FALSE` otherwise.

```
bool_t xdr_setpos(XDR *xdrs, const uint_t pos);
```

A macro that invokes the set position routine associated with the XDR stream *xdrs*. The parameter *pos* is a position value obtained from `xdr_getpos()`. This routine returns `TRUE` if the XDR stream was repositioned, and `FALSE` otherwise.

Warning: it is difficult to reposition some types of XDR streams, so this routine may fail with one type of stream and succeed with another. Therefore, applications written for portability should not depend on this feature.

```
unsigned long xdr_sizeof(xdrproc_t func, void *data);
```

This routine returns the number of bytes required to encode *data* using the XDR filter function *func*, excluding potential overhead such as RPC headers or record markers. `0` is returned on error. This information might be used to select between transport protocols, or to determine the buffer size for various lower levels of RPC client and server creation routines, or to allocate storage when XDR is used outside of the RPC subsystem.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** `malloc(3C)`, `rpc(3NSL)`, `xdr_complex(3NSL)`, `xdr_create(3NSL)`, `xdr_simple(3NSL)`, `attributes(5)`

## xdr\_complex(3NSL)

<b>NAME</b>	xdr_complex, xdr_array, xdr_bytes, xdr_opaque, xdr_pointer, xdr_reference, xdr_string, xdr_union, xdr_vector, xdr_wrapstring – library routines for external data representation
<b>DESCRIPTION</b>	XDR library routines allow C programmers to describe complex data structures in a machine-independent fashion. Protocols such as remote procedure calls (RPC) use these routines to describe the format of the data. These routines are the XDR library routines for complex data structures. They require the creation of XDR streams. See <code>xdr_create(3NSL)</code> .
<b>Routines</b>	<p>See <code>rpc(3NSL)</code> for the definition of the XDR data structure. Note that any buffers passed to the XDR routines must be properly aligned. It is suggested either that <code>malloc()</code> be used to allocate these buffers, or that the programmer insure that the buffer address is divisible evenly by four.</p> <pre>#include &lt;rpc/xdr.h&gt;</pre> <p><code>bool_t xdr_array(XDR *xdrs, caddr_t *arrp, uint_t *sizep, const uint_t maxsize, const uint_t elsize, const xdrproc_t elproc);</code> <code>xdr_array()</code> translates between variable-length arrays and their corresponding external representations. The parameter <code>arrp</code> is the address of the pointer to the array, while <code>sizep</code> is the address of the element count of the array; this element count cannot exceed <code>maxsize</code>. The parameter <code>elsize</code> is the size of each of the array's elements, and <code>elproc</code> is an XDR routine that translates between the array elements' C form and their external representation. If <code>*arrp</code> is NULL when decoding, <code>xdr_array()</code> allocates memory and <code>*arrp</code> points to it. This routine returns TRUE if it succeeds, FALSE otherwise.</p> <p><code>bool_t xdr_bytes(XDR *xdrs, char **sp, uint_t *sizep, const uint_t maxsize);</code> <code>xdr_bytes()</code> translates between counted byte strings and their external representations. The parameter <code>sp</code> is the address of the string pointer. The length of the string is located at address <code>sizep</code>; strings cannot be longer than <code>maxsize</code>. If <code>*sp</code> is NULL when decoding, <code>xdr_bytes()</code> allocates memory and <code>*sp</code> points to it. This routine returns TRUE if it succeeds, FALSE otherwise.</p> <p><code>bool_t xdr_opaque(XDR *xdrs, caddr_t cp, const uint_t cnt);</code> <code>xdr_opaque()</code> translates between fixed size opaque data and its external representation. The parameter <code>cp</code> is the address of the opaque object, and <code>cnt</code> is its size in bytes. This routine returns TRUE if it succeeds, FALSE otherwise.</p> <p><code>bool_t xdr_pointer(XDR *xdrs, char **objpp, uint_t objsize, const xdrproc_t xdrobj);</code> Like <code>xdr_reference()</code> except that it serializes null pointers, whereas <code>xdr_reference()</code> does not. Thus, <code>xdr_pointer()</code> can represent recursive data structures, such as binary trees or linked lists. If <code>*objpp</code> is NULL when decoding, <code>xdr_pointer()</code> allocates memory and <code>*objpp</code> points to it.</p> <p><code>bool_t xdr_reference(XDR *xdrs, caddr_t *pp, uint_t size, const xdrproc_t proc);</code> <code>xdr_reference()</code> provides pointer chasing within structures. The parameter <code>pp</code> is the address of the pointer; <code>size</code> is the <code>sizeof</code> the structure that <code>*pp</code> points to; and <code>proc</code> is an XDR procedure that translates the structure between its C form and</p>

its external representation. If *\*pp* is NULL when decoding, `xdr_reference()` allocates memory and *\*pp* points to it. This routine returns 1 if it succeeds, 0 otherwise.

Warning: this routine does not understand null pointers. Use `xdr_pointer()` instead.

```
bool_t xdr_string(XDR *xdrs, char **sp, const uint_t maxsize);
```

`xdr_string()` translates between C strings and their corresponding external representations. Strings cannot be longer than *maxsize*. Note: *sp* is the address of the string's pointer. If *\*sp* is NULL when decoding, `xdr_string()` allocates memory and *\*sp* points to it. This routine returns TRUE if it succeeds, FALSE otherwise. Note: `xdr_string()` can be used to send an empty string (""), but not a null string.

```
bool_t xdr_union(XDR *xdrs, enum_t *dscmp, char *unp, const struct xdr_discrim
*choices, const xdrproc_t (*defaultarm));
```

`xdr_union()` translates between a discriminated C union and its corresponding external representation. It first translates the discriminant of the union located at *dscmp*. This discriminant is always an `enum_t`. Next the union located at *unp* is translated. The parameter *choices* is a pointer to an array of `xdr_discrim` structures. Each structure contains an ordered pair of [*value*, *proc*]. If the union's discriminant is equal to the associated *value*, then the *proc* is called to translate the union. The end of the `xdr_discrim` structure array is denoted by a routine of value NULL. If the discriminant is not found in the *choices* array, then the *defaultarm* procedure is called (if it is not NULL). It returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_vector(XDR *xdrs, char *arrp, const uint_t size, const uint_t elsize, const
xdrproc_t elproc);
```

`xdr_vector()` translates between fixed-length arrays and their corresponding external representations. The parameter *arrp* is the address of the pointer to the array, while *size* is the element count of the array. The parameter *elsize* is the `sizeof` each of the array's elements, and *elproc* is an XDR routine that translates between the array elements' C form and their external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_wrapstring(XDR *xdrs, char **sp);
```

A routine that calls `xdr_string(xdrs, sp, maxuint)`; where *maxuint* is the maximum value of an unsigned integer.

Many routines, such as `xdr_array()`, `xdr_pointer()`, and `xdr_vector()` take a function pointer of type `xdrproc_t()`, which takes two arguments.

`xdr_string()`, one of the most frequently used routines, requires three arguments, while `xdr_wrapstring()` only requires two. For these routines, `xdr_wrapstring()` is desirable. This routine returns TRUE if it succeeds, FALSE otherwise.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

xdr\_complex(3NSL)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** malloc(3C), rpc(3NSL), xdr\_admin(3NSL), xdr\_create(3NSL), xdr\_simple(3NSL), attributes(5)



<b>NAME</b>	xdr_create, xdr_destroy, xdrmem_create, xdrrec_create, xdrstdio_create – library routines for external data representation stream creation
<b>SYNOPSIS</b>	<pre>#include &lt;rpc/xdr.h&gt;  void <b>xdr_destroy</b>(XDR *xdrs);  void <b>xdrmem_create</b>(XDR *xdrs, const caddr_t addr, const uint_t size,     const enum xdr_op op);  void <b>xdrrec_create</b>(XDR *xdrs, const uint_t sendsz, const uint_t     recvsz, const caddr_t handle, const int (*readit) const void     *read_handle, char *buf, const int len, const int (*writeit) const     void *write_handle, const char *buf, const int len);  void <b>xdrstdio_create</b>(XDR *xdrs, FILE *file, const enum xdr_op op);</pre>
<b>DESCRIPTION</b>	<p>The XDR library routines allow C programmers to describe arbitrary data structures in a machine-independent fashion. Protocols such as remote procedure calls (RPC) use these routines to describe the format of the data.</p> <p>These routines deal with the creation of XDR streams, which must be created before any data can be translated into XDR format.</p>
<b>Routines</b>	<p>See <code>rpc(3NSL)</code> for the definition of the XDR CLIENT and SVCXPRT data structures. Any buffers passed to the XDR routines must be properly aligned. Use <code>malloc(3C)</code> to allocate these buffers or be sure that the buffer address is divisible evenly by four.</p> <p><code>xdr_destroy()</code>  A macro that invokes the destroy routine associated with the XDR stream, <i>xdrs</i>. Private data structures associated with the stream are freed. Using <i>xdrs</i> after <code>xdr_destroy()</code> is invoked is undefined.</p> <p><code>xdrmem_create()</code>  This routine initializes the XDR stream object pointed to by <i>xdrs</i>. The stream's data is written to or read from a chunk of memory at location <i>addr</i> whose length is no less than <i>size</i> bytes long. The <i>op</i> determines the direction of the XDR stream. The value of <i>op</i> can be either XDR_ENCODE, XDR_DECODE, or XDR_FREE.</p> <p><code>xdrrec_create()</code>  This routine initializes the read-oriented XDR stream object pointed to by <i>xdrs</i>. The stream's data is written to a buffer of size <i>sendsz</i>. A value of 0 indicates the system should use a suitable default. The stream's data is read from a buffer of size <i>recvsz</i>. It too can be set to a suitable default by passing a 0 value. When a stream's output buffer is full, <i>writeit</i> is called. Similarly, when a stream's input buffer is empty, <code>xdrrec_create()</code> calls <i>readit</i>. The behavior of these two routines is similar to the system calls <code>read()</code> and <code>write()</code>, except that an appropriate handle, <i>read_handle</i> or <i>write_handle</i>, is passed to the former routines as the first parameter instead of a file descriptor. See <code>read(2)</code> and <code>write(2)</code>, respectively. The XDR stream's <i>op</i> field must be set by the caller.</p>

## xdr\_create(3NSL)

This XDR stream implements an intermediate record stream. Therefore, additional bytes in the stream are provided for record boundary information.

### xdrstdio\_create()

This routine initializes the XDR stream object pointed to by *xdrs*. The XDR stream data is written to or read from the standard I/O stream *file*. The parameter *op* determines the direction of the XDR stream. The value of *op* can be either XDR\_ENCODE, XDR\_DECODE, or XDR\_FREE.

The destroy routine associated with XDR streams calls *fflush()* on the *file* stream, but never *fclose()*. See *fclose(3C)*.

A failure of any of these functions can be detected by first initializing the *x\_ops* field in the XDR structure (*xdrs->x\_ops*) to *NULL* before calling the *xdr\*\_create()* function. If the *x\_ops* field is still *NULL*, after the return from the *xdr\*\_create()* function, the call has failed. If the *x\_ops* field contains some other value, assume that the call has succeeded.

**ATTRIBUTES** See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**SEE ALSO** *read(2)*, *write(2)*, *fclose(3C)*, *malloc(3C)*, *rpc(3NSL)*, *xdr\_admin(3NSL)*, *xdr\_complex(3NSL)*, *xdr\_simple(3NSL)*, *attributes(5)*

<b>NAME</b>	xdr_simple, xdr_bool, xdr_char, xdr_double, xdr_enum, xdr_float, xdr_free, xdr_hyper, xdr_int, xdr_long, xdr_longlong_t, xdr_quadruple, xdr_short, xdr_u_char, xdr_u_hyper, xdr_u_int, xdr_u_long, xdr_u_longlong_t, xdr_u_short, xdr_void – library routines for external data representation
<b>SYNOPSIS</b>	<pre>#include&lt;rpc/xdr.h&gt;  bool_t xdr_bool(XDR *xdrs, bool_t *bp); bool_t xdr_char(XDR *xdrs, char *cp); bool_t xdr_double(XDR *xdrs, double *dp); bool_t xdr_enum(XDR *xdrs, enum_t *ep); bool_t xdr_float(XDR *xdrs, float *fp); void xdr_free(xdrproc_t proc, char *objp); bool_t xdr_hyper(XDR *xdrs, longlong_t *llp); bool_t xdr_int(XDR *xdrs, int *ip); bool_t xdr_long(XDR *xdrs, longt *lp); bool_t xdr_longlong_t(XDR *xdrs, longlong_t *llp); bool_t xdr_quadruple(XDR *xdrs, long double *pq); bool_t xdr_short(XDR *xdrs, short *sp); bool_t xdr_u_char(XDR *xdrs, unsigned char *ucp); bool_t xdr_u_hyper(XDR *xdrs, u_longlong_t *ullp); bool_t xdr_u_int(XDR *xdrs, unsigned *up); bool_t xdr_u_long(XDR *xdrs, unsigned long *ulp); bool_t xdr_u_longlong_t(XDR *xdrs, u_longlong_t *ullp); bool_t xdr_u_short(XDR *xdrs, unsigned short *usp); bool_t xdr_void(void);</pre>
<b>DESCRIPTION</b>	<p>The XDR library routines allow C programmers to describe simple data structures in a machine-independent fashion. Protocols such as remote procedure calls (RPC) use these routines to describe the format of the data.</p> <p>These routines require the creation of XDR streams (see xdr_create(3NSL)).</p>
<b>Routines</b>	<p>See rpc(3NSL) for the definition of the XDR data structure. Note that any buffers passed to the XDR routines must be properly aligned. It is suggested that malloc(3C) be used to allocate these buffers or that the programmer insure that the buffer address is divisible evenly by four.</p> <p>xdr_bool()                      xdr_bool() translates between booleans (C integers) and their external representations. When encoding</p>

## xdr\_simple(3NSL)

	<p>data, this filter produces values of either 1 or 0. This routine returns TRUE if it succeeds, FALSE otherwise.</p>
xdr_char()	<p>xdr_char() translates between C characters and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise. Note: encoded characters are not packed, and occupy 4 bytes each. For arrays of characters, it is worthwhile to consider xdr_bytes(), xdr_opaque(), or xdr_string() (see xdr_complex(3NSL)).</p>
xdr_double()	<p>xdr_double() translates between C double precision numbers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.</p>
xdr_enum()	<p>xdr_enum() translates between C enums (actually integers) and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.</p>
xdr_float()	<p>xdr_float() translates between C floats and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.</p>
xdr_free()	<p>Generic freeing routine. The first argument is the XDR routine for the object being freed. The second argument is a pointer to the object itself. Note: the pointer passed to this routine is not freed, but what it points to is freed (recursively, depending on the XDR routine).</p>
xdr_hyper()	<p>xdr_hyper() translates between ANSI C long long integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.</p>
xdr_int()	<p>xdr_int() translates between C integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.</p>
xdr_long()	<p>xdr_long() translates between C long integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.</p>
	<p>In a 64-bit environment, this routine returns an error if the value of lp is outside the range [INT32_MIN, INT32_MAX]. The xdr_int() routine is recommended in place of this routine.</p>
xdr_longlong_t()	<p>xdr_longlong_t() translates between ANSI C long long integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise. This routine is identical to xdr_hyper().</p>

xdr_quadruple()	xdr_quadruple() translates between IEEE quadruple precision floating point numbers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.
xdr_short()	xdr_short() translates between C short integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.
xdr_u_char()	xdr_u_char() translates between unsigned C characters and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.
xdr_u_hyper()	xdr_u_hyper() translates between unsigned ANSI C long long integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.
xdr_u_int()	A filter primitive that translates between a C unsigned integer and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.
xdr_u_long()	xdr_u_long() translates between C unsigned long integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.  In a 64-bit environment, this routine returns an error if the value of <i>ulp</i> is outside the range [0, <code>UINT32_MAX</code> ]. The <code>xdr_u_int()</code> routine is recommended in place of this routine.
xdr_u_longlong_t()	xdr_u_longlong_t() translates between unsigned ANSI C long long integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise. This routine is identical to <code>xdr_u_hyper()</code> .
xdr_u_short()	xdr_u_short() translates between C unsigned short integers and their external representations. This routine returns TRUE if it succeeds, FALSE otherwise.
xdr_void()	This routine always returns TRUE. It may be passed to RPC routines that require a function parameter, where nothing is to be done.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

xdr\_simple(3NSL)

**SEE ALSO** malloc(3C), rpc(3NSL), xdr\_admin(3NSL), xdr\_complex(3NSL),  
xdr\_create(3NSL), attributes(5)

<b>NAME</b>	xfn – overview of the XFN interface
<b>DESCRIPTION</b>	<p>The primary service provided by a federated naming system is to map a <i>composite name</i> to a <i>reference</i>. A composite name is composed of name components from one or more naming systems. A reference consists of one or more communication end points. An additional service provided by a federated naming system is to provide access to attributes associated with named objects. This extension is to satisfy most applications' additional naming service needs without cluttering the basic naming service model. XFN is a programming interface for a federated naming service.</p> <p>To use the XFN interface, include the <code>xfn/xfn.h</code> header file and link the application with <code>-lxfn</code>.</p> <p>The <code>xfn/xfn.h</code> header file contains the interface declarations for:</p> <ul style="list-style-type: none"> <li>■ the XFN base context interface,</li> <li>■ the XFN base attribute interface,</li> <li>■ status object and status codes used by operations in these two interfaces,</li> <li>■ abstract data types passed as parameters to and returned as values from operations in these two interfaces, and</li> <li>■ the interface for the XFN standard syntax model for parsing compound names.</li> </ul>
<b>FILES</b>	<code>/usr/include/xfn/xfn.h</code>
<b>SEE ALSO</b>	<code>FN_ctx_t(3XFN)</code> , <code>FN_status_t(3XFN)</code> , <code>xfn_attributes(3XFN)</code> , <code>xfn_composite_names(3XFN)</code> , <code>xfn_compound_names(3XFN)</code> , <code>xfn_status_codes(3XFN)</code> , <code>fns(5)</code> , <code>fns_policies(5)</code>
<b>NOTES</b>	The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## xfn\_attributes(3XFN)

<b>NAME</b>	xfn_attributes – an overview of XFN attribute operations
<b>DESCRIPTION</b>	<p>XFN assumes the following model for attributes. A set of zero or more attributes is associated with a named object. Each attribute in the set has a unique attribute identifier, an attribute syntax, and a (possibly empty) set of distinct data values. Each attribute value has an opaque data type. The attribute identifier serves as a name for the attribute. The attribute syntax indicates how the value is encoded in the buffer.</p> <p>The operations of the base attribute interface may be used to examine and modify the settings of attributes associated with existing named objects. These objects may be contexts or other types of objects. The attribute operations do not create names or remove names from contexts.</p> <p>The range of support for attribute operations may vary widely. Some naming systems may not support any attribute operations. Other naming systems may only support read operations, or operations on attributes whose identifiers are in some fixed set. A naming system may limit attributes to have a single value, or may require at least one value. Some naming systems may only associate attributes with context objects, while others may allow associating attributes with non-context objects.</p> <p>These are the interfaces:</p> <pre>#include &lt;xfn/xfn.h&gt;  FN_attribute_t *fn_attr_get(FN_ctx_t *ctx, const FN_composite_name_t *name,     const FN_identifier_t *attribute_id, FN_status_t *status);  int fn_attr_modify(FN_ctx_t *ctx, const FN_composite_name_t *name,     unsigned int mod_op, const FN_attribute_t *attr, FN_status_t *status);  FN_attrset_t *fn_attr_get_ids(FN_ctx_t *ctx, const FN_composite_name_t *name,     FN_status_t *status);  FN_valuelist_t *fn_attr_get_values(FN_ctx_t *ctx,     const FN_composite_name_t *name,     const FN_identifier_t *attribute_id, FN_status_t *status);  FN_attrvalue_t *fn_valuelist_next(FN_valuelist_t *vl,     FN_identifier_t **attr_syntax,     FN_status_t *status);  void fn_valuelist_destroy(FN_valuelist_t *vl, FN_status_t *status);  FN_multigetlist_t *fn_attr_multi_get(FN_ctx_t *ctx,     const FN_composite_name_t *name, const FN_attrset_t *attr_ids,     FN_status_t *status);  FN_attribute_t *fn_multigetlist_next(FN_multigetlist_t *ml,</pre>



```

    FN_status_t *status);

void fn_multigetlist_destroy(FN_multigetlist_t *ml, FN_status_t *status);

int fn_attr_multi_modify(FN_ctx_t *ctx, const FN_composite_name_t *name,
    const FN_attrmodlist_t *mods, FN_status_t *status,
    FN_attrmodlist_t **unexecuted_mods);

FN_attrset_t *fn_ctx_get_syntax_attrs(FN_ctx_t *ctx,
    const FN_composite_name_t *name, FN_status_t *status);

```

The following describes briefly the operations in the base attribute interface. Detailed descriptions are given in the respective reference manual pages for these operations.

`fn_attr_get()` returns the attribute identified. `fn_attr_modify()` modifies the attribute identified as described by `mod_op`.

`fn_attr_get_ids()` returns the identifiers of the attributes of the named object.

`fn_attr_get_values()` and its set of related operations are used for returning the individual values of an attribute.

`fn_attr_multi_get()` and its set of related operations are used for returning the requested attributes associated with the named object. `fn_attr_multi_modify()` modifies multiple attributes associated with the named object in a single invocation.

`fn_ctx_get_syntax_attrs()` returns the syntax attributes associated with the named context.

**ERRORS**

*status* is set as described in `FN_status_t(3XFN)` and `xfn_status_codes(3XFN)`. The following status codes are of special relevance to attribute operations:

<code>FN_E_ATTR_VALUE_REQUIRED</code>	The operation attempted to create an attribute without a value, and the specific naming system does not allow this.
<code>FN_E_ATTR_NO_PERMISSION</code>	The caller did not have permission to perform the attempted attribute operation.
<code>FN_E_INSUFFICIENT_RESOURCES</code>	There are insufficient resources to retrieve the requested attribute(s).
<code>FN_E_INVALID_ATTR_IDENTIFIER</code>	The attribute identifier was not in a format acceptable to the naming system, or its contents was not valid for the format specified for the identifier.
<code>FN_E_INVALID_ATTR_VALUE</code>	One of the values supplied was not in the appropriate form for the given attribute.

## xfn\_attributes(3XFN)

	FN_E_NO_SUCH_ATTRIBUTE	The object did not have an attribute with the given identifier.
	FN_E_TOO_MANY_ATTR_VALUES	The operation attempted to associate more values with an attribute than the naming system supported.
<b>USAGE</b>	<p>Except for <code>fn_ctx_get_syntax_attrs()</code>, an attribute operation using a composite name is not necessarily equivalent to an independent <code>fn_ctx_lookup()</code> operation followed by an attribute operation in which the caller supplies the resulting reference and an empty name. This is because there is a range of attribute models in which an attribute is associated with a name in a context, or an attribute is associated with the object named, or both. XFN accommodates all of these alternatives. Invoking an attribute operation using the target context and the terminal atomic name accesses either the attributes that are associated with the target name or target named object; this is dependent on the underlying attribute model. This document uses the term <i>attributes associated with a named object</i> to refer to all of these cases.</p> <p>XFN specifies no guarantees about the relationship between the attributes and the reference associated with a given name. Some naming systems may store the reference bound to a name in one or more attributes associated with a name. Attribute operations might affect the information used to construct a reference.</p> <p>To avoid undefined results, programmers must use the operations in the context interface and not attribute operations when the intention is to manipulate a reference. Programmers should avoid the use of specific knowledge about how an XFN context implementation over a particular naming system constructs references.</p>	
<b>SEE ALSO</b>	FN_attribute_t(3XFN), FN_attrset_t(3XFN), FN_attrvalue_t(3XFN), FN_composite_name_t(3XFN), FN_ctx_t(3XFN), FN_identifier_t(3XFN), FN_status_t(3XFN), fn_attr_get(3XFN), fn_attr_get_ids(3XFN), fn_attr_get_values(3XFN), fn_attr_modify(3XFN), fn_attr_multi_get(3XFN), fn_attr_multi_modify(3XFN), fn_ctx_get_syntax_attrs(3XFN), fn_ctx_lookup(3XFN), xfn(3XFN), xfn_status_codes(3XFN)	
<b>NOTES</b>	<p>The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.</p>	

<b>NAME</b>	xfn_composite_names – XFN composite syntax: an overview of the syntax for XFN composite name
<b>DESCRIPTION</b>	<p>An <i>XFN composite name</i> consists of an ordered list of zero or more components. Each component is a string name from the namespace of a single naming system. It may be an atomic or a compound name in that namespace.</p> <p>XFN defines an abstract data type, <code>FN_composite_name_t</code>, for representing the structural form of a composite name. XFN also defines a standard string form for composite names. This form is the concatenation of the components of a composite name from left to right with the <i>XFN component separator</i> ('/') character to separate each component.</p> <p>These are the interfaces:</p> <pre data-bbox="456 737 1370 814">#include &lt;xfn/xfn.h&gt; FN_composite_name_t *fn_composite_name_from_string( const FN_string_t *str); FN_string_t *fn_string_from_composite_name( const FN_composite_name_t *name);</pre> <p>The function <code>fn_composite_name_from_string</code> parses the string representation of a composite name into its corresponding composite name object <code>FN_composite_name_t</code>. The function <code>fn_string_from_composite_name</code> composes the string representation of a composite name given its composite name object form <code>FN_composite_name_t</code>.</p>
<b>APPLICATION USAGE</b>	<p>Special characters used in the XFN composite name syntax, such as the separator or escape characters, have the same encoding as they would in ISO 646.</p> <p>All XFN implementations are required to support the portable representation, ISO 646. All other representations are optional.</p> <p>All characters of the string form of a XFN composite name use a single encoding. This does not preclude component names of a composite name in its structural form from having different encodings. Code set mismatches that occur during the process of converting a composite name structure to its string form are resolved in an implementation-dependent way. When an implementation discovers that a composite name has components with incompatible code sets, it returns the error code <code>FN_E_INCOMPATIBLE_CODE_SETS</code>.</p>
<b>SEE ALSO</b>	FN_string_t(3XFN), FN_compound_name_t(3XFN), xfn(3XFN)

## xfn\_compound\_names(3XFN)

<b>NAME</b>	xfn_compound_names – XFN compound syntax: an overview of XFN model for compound name parsing
<b>DESCRIPTION</b>	<p>Each naming system in an XFN federation has a naming convention. XFN defines a standard model of expressing compound name syntax that covers a large number of specific name syntaxes and is expressed in terms of syntax properties of the naming convention.</p> <p>The model uses the attributes in the following table to describe properties of the syntax. Unless otherwise qualified, these syntax attributes have attribute identifiers that use the FN_ID_STRING format. A context that supports the XFN standard syntax model has an attribute set containing the fn_syntax_type (with identifier format FN_ID_STRING) attribute with the value "standard" (ASCII attribute syntax).</p> <p>These are the interfaces:</p> <pre>#include &lt;xfn/xfn.h&gt; FN_attrset_t *fn_ctx_get_syntax_attrs(FN_ctx_t *ctx, const FN_composite_name_t *name, FN_status_t *status); FN_compound_name_t *fn_compound_name_from_syntax_attrs(const FN_attrset_t *aset, const FN_string_t *name, FN_status_t *status);</pre> <p><b>fn_syntax_type</b> Its value is the ASCII string "standard" if the context supports the XFN standard syntax model. Its value is an implementation-specific value if another syntax model is supported.</p> <p><b>fn_std_syntax_direction</b> Its value is an ASCII string, one of "left_to_right", "right_to_left", or "flat". This determines whether the order of components in a compound name string goes from left to right, right to left, or whether the namespace is flat (in other words, not hierarchical; em all names are atomic).</p> <p><b>fn_std_syntax_separator</b> Its value is the separator string for this name syntax. This attribute is required unless the fn_std_syntax_direction is "flat".</p> <p><b>fn_std_syntax_escape</b> If present, its value is the escape string for this name syntax.</p> <p><b>fn_std_syntax_case_insensitive</b> If this attribute is present, it indicates that names that differ only in case are considered identical. If this attribute is absent, it indicates that case is significant. If a value is present, it is ignored.</p> <p><b>fn_std_syntax_begin_quote</b> If present, its value is the begin-quote string for this syntax. There can be multiple values for this attribute.</p>

`fn_std_syntax_end_quote`

If present, its value is the end-quote string for this syntax. There can be multiple values for this attribute.

`fn_std_syntax_ava_separator`

If present, its value is the attribute value assertion separator string for this syntax.

`fn_std_syntax_typeval_separator`

If present, its value is the attribute type-value separator string for this syntax.

`fn_std_syntax_code_sets`

If present, its value identifies the code sets of the string representation for this syntax. Its value consists of a structure containing an array of code sets supported by the context; the first member of the array is the preferred code set of the context. The values for the code sets are defined in the X/Open code set registry. If this attribute is not present, or if the value is empty, the default code set is ISO 646 (same encoding as ASCII).

`fn_std_syntax_locale_info`

If present, identifies locale information, such as character set information, of the string representation for this syntax. The interpretation of its value is implementation-dependent.

The XFN standard syntax attributes are interpreted according to the following rules:

1. In a string without quotes or escapes, any instance of the separator string delimits two atomic names.
2. A separator, quotation or escape string is escaped if preceded immediately (on the left) by the escape string.
3. A non-escaped begin-quote which precedes a component must be matched by a non-escaped end-quote at the end of the component. Quotes embedded in non-quoted names are treated as simple characters and do not need to be matched. An unmatched quotation fails with the status code `FN_E_ILLEGAL_NAME`.
4. If there are multiple values for begin-quote and end-quote, a specific begin-quote value must be matched with its corresponding end-quote value.
5. When the separator appears between a (non-escaped) begin quote and the end quote, it is ignored.
6. When the separator is escaped, it is ignored. An escaped begin-quote or end-quote string is not treated as a quotation mark. An escaped escape string is not treated as an escape string.
7. A non-escaped escape string appearing within quotes is interpreted as an escape string. This can be used to embed an end-quote within a quoted string.

After constructing a compound name from a string, the resulting component atoms have one level of escape strings and quotations interpreted and consumed.

`fn_ctx_get_syntax_attrs()` is used to obtain the syntax attributes associated with a context.

## xfn\_compound\_names(3XFN)

	<code>fn_compound_name_from_syntax()</code> is used to construct a compound name object using the string form of the name and the syntax attributes of the name.
<b>ERRORS</b>	
	<code>FN_E_ILLEGAL_NAME</code> The name supplied to the operation was not a well-formed component according to the name syntax of the context.
	<code>FN_E_INCOMPATIBLE_CODE_SETS</code> Code set mismatches that occur during the construction of the compound name's string form are resolved in an implementation-dependent way. When an implementation discovers that a compound name has components with incompatible code sets, it returns this error code.
	<code>FN_E_INVALID_SYNTAX_ATTRS</code> The syntax attributes supplied are invalid or insufficient to fully specify the syntax.
	<code>FN_E_SYNTAX_NOT_SUPPORTED</code> The syntax specified is not supported.
<b>USAGE</b>	Most applications treat names as opaque data. Hence, the majority of clients of the XFN interface will not need to parse compound names from specific naming systems. Some applications, however, such as browsers, need such capabilities. These applications would use <code>fn_ctx_get_syntax_attrs()</code> to obtain the syntax-related attributes of a context and, if the context uses the XFN standard syntax model, it would examine these attributes to determine the name syntax of the context.
<b>SEE ALSO</b>	<code>FN_attribute_t(3XFN)</code> , <code>FN_attrset_t(3XFN)</code> , <code>FN_compound_name_t(3XFN)</code> , <code>FN_identifier_t(3XFN)</code> , <code>FN_string_t(3XFN)</code> <code>fn_ctx_get_syntax_attrs(3XFN)</code> , <code>xfn(3XFN)</code>
<b>NOTES</b>	The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

<b>NAME</b>	xfn_links – XFN links: an overview of XFN links
<b>DESCRIPTION</b>	<p>An <i>XFN link</i> is a special form of reference that contains a composite name, the <i>link name</i>, and that may be bound to an atomic name in an XFN context. Because the link name is a composite name, it may span multiple namespaces.</p> <p>Normal resolution of names in context operations always follows XFN links. If the first composite name component of the link name is the atomic name ".", the link name is resolved relative to the same context in which the link is bound, otherwise, the link name is resolved relative to the XFN Initial Context of the client. The link name may itself cause resolution to pass through other XFN links. This gives rise to the possibility of a cycle of links whose resolution could not terminate normally. As a simple means to avoid such non-terminating resolutions, implementations may define limits on the number of XFN links that may be resolved in any single operation invoked by the caller.</p> <p>These are the interfaces:</p> <pre> #include &lt;xfn/xfn.h&gt;  FN_ref_t *fn_ref_create_link(const FN_composite_name_t *link_name);  int fn_ref_is_link(const FN_ref_t *ref);  FN_composite_name_t *fn_ref_link_name(const FN_ref_t *link_ref);  FN_ref_t *fn_ctx_lookup_link(FN_ctx_t *ctx, const FN_composite_name_t *name,                              FN_status_t *status);  unsigned int fn_status_link_code(const FN_status_t *stat);  const FN_composite_name_t *fn_status_link_remaining_name(     const FN_status_t *stat);  const FN_composite_name_t *fn_status_link_resolved_name(     const FN_status_t *stat);  const FN_ref_t *fn_status_link_resolved_ref(const FN_status_t *stat);  int fn_status_set_link_code(FN_status_t *stat,                            unsigned int code);  int fn_status_set_link_remaining_name(FN_status_t *stat,                                      const FN_composite_name_t *name);  int fn_status_set_link_resolved_name(FN_status_t *stat,                                      const FN_composite_name_t *name);  int fn_status_set_link_resolved_ref(FN_status_t *stat,                                     const FN_ref_t *ref); </pre>

## xfn\_links(3XFN)

Links are bound to names using the normal `fn_ctx_bind()` and unbound using the normal `fn_ctx_unbind()` operation. The operation `fn_ref_create_link()` is provided for constructing a link reference from a composite name. Since normal resolution always follows links, a separate operation, `fn_ctx_lookup_link()` is provided to lookup the link itself.

In the case that an error occurred while resolving an XFN link, the status object set by the operation contains additional information about that error and sets the corresponding link status fields using `fn_status_set_link_code()`, `fn_status_set_link_remaining_name()`, `fn_status_set_link_resolved_name()` and `fn_status_set_link_resolved_ref()`. The link status fields can be retrieved using `fn_status_link_code()`, `fn_status_link_remaining_name()`, `fn_status_link_resolved_name()` and `fn_status_link_resolved_ref()`.

### ERRORS

The following status codes are of special relevance when performing operations involving XFN links:

<code>FN_E_LINK_ERROR</code>	There was an error encountered resolving an XFN link encountered during resolution of the supplied name. Check the link part of the status object to determine cause of the link error.
<code>FN_E_LINK_LOOP_LIMIT</code>	A non-terminating loop (cycle) in the resolution can arise due to XFN links encountered during the resolution of a composite name. This code indicates either the definite detection of such a cycle, or that resolution exceeded an implementation-defined limit on the number of XFN links allowed for a single operation invoked by the caller.
<code>FN_E_MALFORMED_LINK</code>	A malformed link reference was encountered. For the <code>fn_ctx_lookup_link()</code> operation, the name supplied resolved to a reference that was not a link.

### APPLICATION USAGE

For the `fn_ctx_bind()`, `fn_ctx_unbind()`, `fn_ctx_rename()`, `fn_ctx_lookup_link()`, `fn_ctx_create_subcontext()` and `fn_ctx_destroy_subcontext()` operations, resolution of the given name continues to the target context — that named by all but the terminal atomic part of the given name; the terminal atomic name is not resolved. Consequently, for operations that involve unbinding the terminal atomic part such as `fn_ctx_unbind()`, if the terminal atomic name is bound to a link, the link is not followed and the link itself is unbound from the terminal atomic name.

Many naming systems support a native notion of link that may be used within the naming system itself. XFN does not determine whether there is any relationship between such native links and XFN links.



xfn\_links(3XFN)

**SEE ALSO** FN\_composite\_name\_t(3XFN), FN\_ref\_t(3XFN), FN\_status\_t(3XFN),  
fn\_ctx\_bind(3XFN), fn\_ctx\_destroy\_subcontext(3XFN),  
fn\_ctx\_lookup(3XFN), fn\_ctx\_lookup\_link(3XFN), fn\_ctx\_rename(3XFN),  
fn\_ctx\_unbind(3XFN), xfn\_status\_codes(3XFN), xfn(3XFN)

## xfn\_status\_codes(3XFN)

<b>NAME</b>	xfn_status_codes – descriptions of XFN status codes																		
<b>SYNOPSIS</b>	<pre>#include &lt;xfn/xfn.h&gt;</pre>																		
<b>DESCRIPTION</b>	<p>The result status of operations in the context interface and the attribute interface is encapsulated in an <code>FN_status_t</code> object. This object contains information about how the operation completed: whether an error occurred in performing the operation; if so, what kind of error; and information localizing where the error occurred. In the case that the error occurred while resolving an XFN link, the status object contains additional information about that error.</p> <p>The context status object consists of several items of information. One of them is the primary status code, describing the disposition of the operation. In the case that an error occurred while resolving an XFN link, the primary status code has the value <code>FN_E_LINK_ERROR</code>, and the link status code describes the error that occurred while resolving the XFN link.</p>																		
<b>XFN Status Codes</b>	<p>Both the primary status code and the link status code are values of type <code>unsigned int</code> that are drawn from the same set of meaningful values. XFN reserves the values 0 through 127 for standard meanings. Currently, values and interpretations for the following codes are determined by XFN.</p> <table><tr><td><code>FN_SUCCESS</code></td><td>The operation succeeded.</td></tr><tr><td><code>FN_E_ATTR_NO_PERMISSION</code></td><td>The caller did not have permission to perform the attempted attribute operation.</td></tr><tr><td><code>FN_E_ATTR_VALUE_REQUIRED</code></td><td>The operation attempted to create an attribute without a value, and the specific naming system does not allow this.</td></tr><tr><td><code>FN_E_AUTHENTICATION_FAILURE</code></td><td>The identity of the client principal could not be verified.</td></tr><tr><td><code>FN_E_COMMUNICATION_FAILURE</code></td><td>An error occurred in communicating with one of the contexts involved in the operation.</td></tr><tr><td><code>FN_E_CONFIGURATION_ERROR</code></td><td>A problem was detected that indicated an error in the installation of the XFN implementation.</td></tr><tr><td><code>FN_E_CONTINUE</code></td><td>The operation should be continued using the remaining name and the resolved reference returned in the status.</td></tr><tr><td><code>FN_E_CTX_NO_PERMISSION</code></td><td>The client did not have permission to perform the operation.</td></tr><tr><td><code>FN_E_CTX_NOT_EMPTY</code></td><td>(Applies only to <code>fn_ctx_destroy_subcontext()</code>.) The</td></tr></table>	<code>FN_SUCCESS</code>	The operation succeeded.	<code>FN_E_ATTR_NO_PERMISSION</code>	The caller did not have permission to perform the attempted attribute operation.	<code>FN_E_ATTR_VALUE_REQUIRED</code>	The operation attempted to create an attribute without a value, and the specific naming system does not allow this.	<code>FN_E_AUTHENTICATION_FAILURE</code>	The identity of the client principal could not be verified.	<code>FN_E_COMMUNICATION_FAILURE</code>	An error occurred in communicating with one of the contexts involved in the operation.	<code>FN_E_CONFIGURATION_ERROR</code>	A problem was detected that indicated an error in the installation of the XFN implementation.	<code>FN_E_CONTINUE</code>	The operation should be continued using the remaining name and the resolved reference returned in the status.	<code>FN_E_CTX_NO_PERMISSION</code>	The client did not have permission to perform the operation.	<code>FN_E_CTX_NOT_EMPTY</code>	(Applies only to <code>fn_ctx_destroy_subcontext()</code> .) The
<code>FN_SUCCESS</code>	The operation succeeded.																		
<code>FN_E_ATTR_NO_PERMISSION</code>	The caller did not have permission to perform the attempted attribute operation.																		
<code>FN_E_ATTR_VALUE_REQUIRED</code>	The operation attempted to create an attribute without a value, and the specific naming system does not allow this.																		
<code>FN_E_AUTHENTICATION_FAILURE</code>	The identity of the client principal could not be verified.																		
<code>FN_E_COMMUNICATION_FAILURE</code>	An error occurred in communicating with one of the contexts involved in the operation.																		
<code>FN_E_CONFIGURATION_ERROR</code>	A problem was detected that indicated an error in the installation of the XFN implementation.																		
<code>FN_E_CONTINUE</code>	The operation should be continued using the remaining name and the resolved reference returned in the status.																		
<code>FN_E_CTX_NO_PERMISSION</code>	The client did not have permission to perform the operation.																		
<code>FN_E_CTX_NOT_EMPTY</code>	(Applies only to <code>fn_ctx_destroy_subcontext()</code> .) The																		

xfn\_status\_codes(3XFN)

	naming system required that the context be empty before its destruction, and it was not empty.
FN_E_CTX_UNAVAILABLE	Service could not be obtained from one of the contexts involved in the operation. This may be because the naming system is busy, or is not providing service. In some implementations this may not be distinguished from a communication failure.
FN_E_ILLEGAL_NAME	The name supplied to the operation was not a well- formed XFN composite name, or one of the component names was not well-formed according to the syntax of the naming system(s) involved in its resolution.
FN_E_E_INCOMPATIBLE_CODE_SETS	The operation involved character strings of incompatible code sets, or the supplied code set is not supported by the implementation.
FN_E_INSUFFICIENT_RESOURCES	Either the client or one of the involved contexts could not obtain sufficient resources (for example, memory, file descriptors, communication ports, stable media space, and so on) to complete the operation successfully.
FN_E_INVALID_ATTR_IDENTIFIER	The attribute identifier was not in a format acceptable to the naming system, or its content was not valid for the format specified for the identifier.
FN_E_INVALID_ATTR_VALUE	One of the values supplied was not in the appropriate form for the given attribute.
FN_E_INVALID_ENUM_HANDLE	The enumeration handle supplied was invalid, either because it was from another enumeration, or because an update operation occurred during the enumeration, or because of some other reason.
FN_E_INVALID_SYNTAX_ATTRS	The syntax attributes supplied are invalid or insufficient to fully specify the syntax.
FN_E_LINK_ERROR	There was an error in resolving an XFN link encountered during resolution of the supplied name.
FN_E_LINK_LOOP_LIMIT	A non-terminating loop (cycle) in the resolution can arise due to XFN links

## xfn\_status\_codes(3XFN)

	encountered during the resolution of a composite name. This code indicates either the definite detection of such a cycle, or that resolution exceeded an implementation-defined limit on the number of XFN links allowed for a single operation invoked by the caller.
FN_E_MALFORMED_LINK	A malformed link reference was encountered. For <code>fn_ctx_lookup_link()</code> , the name supplied resolved to a reference that was not a link.
FN_E_MALFORMED_REFERENCE	A context object could not be constructed from the supplied reference, because the reference was not properly formed.
FN_E_NAME_IN_USE	(Only for operations that bind names.) The supplied name was already in use.
FN_E_NAME_NOT_FOUND	Resolution of the supplied composite name proceeded to a context in which the next atomic component of the name was not bound.
FN_E_NO_SUCH_ATTRIBUTE	The object did not have an attribute with the given identifier.
FN_E_NO_SUPPORTED_ADDRESS	A context object could not be constructed from a particular reference. The reference contained no address type over which the context interface was supported.
FN_E_NOT_A_CONTEXT	Either one of the intermediate atomic names did not name a context, and resolution could not proceed beyond this point, or the operation required that the caller supply the name of a context, and the name did not resolve to a reference for a context.
FN_E_OPERATION_NOT_SUPPORTED	The operation attempted is not supported.
FN_E_PARTIAL_RESULT	The operation attempted is returning a partial result.
FN_E_SYNTAX_NOT_SUPPORTED	The syntax type specified is not supported.
FN_E_TOO_MANY_ATTR_VALUES	The operation attempted to associate more values with an attribute than the naming system supported.

xfn\_status\_codes(3XFN)

FN\_E\_UNSPECIFIED\_ERROR      An error occurred that could not be classified by any of the other error codes.

**FILES**      #include <xfn/xfn.h>      XFN status codes header file

**SEE ALSO**      FN\_status\_t(3XFN), xfn(3XFN)

**NOTES**      The implementation of XFN in this Solaris release is based on the X/Open preliminary specification. It is likely that there will be minor changes to these interfaces to reflect changes in the final version of this specification. The next minor release of Solaris will offer binary compatibility for applications developed using the current interfaces. As the interfaces evolve toward standardization, it is possible that future releases of Solaris will require minor source code changes to applications that have been developed against the preliminary specification.

## ypclnt(3NSL)

<b>NAME</b>	<code>ypclnt</code> , <code>yp_get_default_domain</code> , <code>yp_bind</code> , <code>yp_unbind</code> , <code>yp_match</code> , <code>yp_first</code> , <code>yp_next</code> , <code>yp_all</code> , <code>yp_order</code> , <code>yp_master</code> , <code>yperr_string</code> , <code>ypprot_err</code> – NIS Version 2 client interface
<b>SYNOPSIS</b>	<pre>cc -flag ... file ...-lnsl [library ...] #include &lt;rpcsvc/ypclnt.h&gt; #include &lt;rpcsvc/yp_prot.h&gt;  <b>yp_bind</b>(char *indomain);  void <b>yp_unbind</b> (char *indomain);  <b>yp_get_default_domain</b>(char **outdomain);  <b>yp_match</b>(char *indomain, char *inmap, char *inkey, int inkeylen, char **outval, int *outvallen);  <b>yp_first</b>(char *indomain, char *inmap, char **outkey, int *outkeylen, char **outval, int *outvallen);  <b>yp_next</b>(char *indomain, char *inmap, char *inkey, int *inkeylen, char **outkey, int *outkeylen, char **outval, int *outvallen);  <b>yp_all</b>(char *indomain, char *inmap, struct ypall_callback *incallback);  <b>yp_order</b>(char *indomain, char *inmap, unsigned long *outorder);  <b>yp_master</b>(char *indomain, char *inmap, char **outname);  char *<b>yperr_string</b>(int incode);  <b>ypprot_err</b>(unsigned int incode);</pre>
<b>DESCRIPTION</b>	<p>This package of functions provides an interface to NIS, Network Information Service Version 2, formerly referred to as YP. In this version of SunOS, NIS version 2 is supported only for compatibility with previous versions. The recommended enterprise level information service is NIS+ or NIS version 3. See <code>nis+(1)</code>. The current SunOS supports only the client interface to NIS version 2. This client interface will in turn be served either by an existing <code>ypserv</code> process running on another machine on the network that has an earlier version of SunOS, or by a NIS+ server running in "YP-compatibility mode." See <code>rpc.nisd(1M)</code>. The NOTES section in <code>ypfiles(4)</code> discusses the implications of being an NIS client of an NIS+ server in "YP-compatibility mode." For commands to access NIS from a client machine, see <code>ypbind(1M)</code>, <code>ypwhich(1)</code>, <code>ypmatch(1)</code>, and <code>ypcat(1)</code>. The package can be loaded from the standard library, <code>/usr/lib/libnsl.so.1</code>.</p> <p>All input parameter names begin with <i>in</i>. Output parameters begin with <i>out</i>. Output parameters of type <code>char **</code> should be addresses of uninitialized character pointers. Memory is allocated by the NIS client package using <code>malloc(3C)</code> and can be freed by the user code if it has no continuing need for it. For each <i>outkey</i> and <i>outval</i>, two extra bytes of memory are allocated at the end that contain NEWLINE and null, respectively, but these two bytes are not reflected in <i>outkeylen</i> or <i>outvallen</i>. The</p>

*indomain* and *inmap* strings must be non-null and null-terminated. String parameters that are accompanied by a count parameter may not be null, but they may point to null strings, with the count parameter indicating this. Counted strings need not be null-terminated.

All functions in this package of type *int* return 0 if they succeed. Otherwise, they return a failure code (YPERR\_XXX). Failure codes are described in the ERRORS section.

## Routines

`yp_bind()`

To use the NIS name services, the client process must be “bound” to an NIS server that serves the appropriate domain using `yp_bind()`. Binding need not be done explicitly by user code. Binding is done automatically whenever an NIS lookup function is called. `yp_bind()` can be called directly for processes that make use of a backup strategy, for example, a local file in cases when NIS services are not available. If a process calls `yp_bind()`, it should call `yp_unbind()` when it is finished using NIS in order to free up resources.

`yp_unbind()`

Each binding allocates or uses up one client process socket descriptor. Each bound domain costs one socket descriptor. However, multiple requests to the same domain use that same descriptor. `yp_unbind()` is available at the client interface for processes that explicitly manage their socket descriptors while accessing multiple domains. The call to `yp_unbind()` makes the domain *unbound*, and frees all per-process and per-node resources used to bind it.

If an RPC failure results upon use of a binding, that domain will be unbound automatically. At that point, the `ypclnt()` layer will retry a few more times or until the operation succeeds, provided that `rpcbind(1M)` and `ypbind(1M)` are running, and either:

- The client process cannot bind a server for the proper domain; or
- RPC requests to the server fail.

Under the following circumstances, the `ypclnt` layer will return control to the user code, with either an error or success code and the results:

- If an error is not RPC-related.
- If `rpcbind` is not running.
- If `ypbind` is not running.
- If a bound `ypserv` process returns any answer (success or failure).

`yp_get_default_domain()`

The NIS lookup calls require a map name and a domain name, at minimum. The client process should know the name of the map of interest. Client processes fetch the node’s default domain by calling `yp_get_default_domain()` and use the returned *outdomain* as the *indomain* parameter to successive NIS name service calls. The domain returned is the same as that returned using the `SI_SRPC_DOMAIN` command to the `sysinfo(2)` system call. The value returned in *outdomain* should not be freed.

## ypclnt(3NSL)

### yp\_match()

`yp_match()` returns the value associated with a passed key. This key must be exact because no pattern matching is available. `yp_match()` requires a full YP map name, such as `hosts.byname`, instead of the nickname `hosts`.

### yp\_first()

`yp_first()` returns the first key-value pair from the named map in the named domain.

### yp\_next()

`yp_next()` returns the next key-value pair in a named map. The *inkey* parameter must be the *outkey* returned from an initial call to `yp_first()` (to get the second key-value pair) or the one returned from the *n*th call to `yp_next()` (to get the *n*th + second key-value pair). Similarly, the *inkeylen* parameter must be the *outkeylen* returned from the earlier `yp_first()` or `yp_next()` call.

The concept of first and next is particular to the structure of the NIS map being processed. Retrieval order is not related to either the lexical order within any original (non-NIS name service) data base, or to any obvious numerical sorting order on the keys, values, or key-value pairs. The only ordering guarantee is that if the `yp_first()` function is called on a particular map, and then the `yp_next()` function is repeatedly called on the same map at the same server until the call fails with a reason of `YPERR_NOMORE`, every entry in the data base is seen exactly once. Further, if the same sequence of operations is performed on the same map at the same server, the entries are seen in the same order.

Under conditions of heavy server load or server failure, the domain can become unbound, then bound once again (perhaps to a different server) while a client is running. This binding can cause a break in one of the enumeration rules. Specific entries may be seen twice by the client, or not at all. This approach protects the client from error messages that would otherwise be returned in the midst of the enumeration. For a better solution to enumerating all entries in a map, see `yp_all()`.

### yp\_all()

The function `yp_all()` provides a way to transfer an entire map from server to client in a single request using TCP (rather than UDP as with other functions in this package). The entire transaction takes place as a single RPC request and response. `yp_all()` can be used just like any other NIS name service procedure to identify the map in the normal manner and to supply the name of a function that will be called to process each key-value pair within the map. The call to `yp_all()` returns only when the transaction is completed (successfully or unsuccessfully), or the `foreach()` function decides that it does not want to see any more key-value pairs.

The third parameter to `yp_all()` is:

```
struct ypall_callback *incallback {
    int (*foreach)();
    char *data;
};
```



The function `foreach()` is called:

```
foreach(int instatus, char *inkey,
int inkeylen, char *inval,
int invallen, char *indata);
```

The *instatus* parameter holds one of the return status values defined in `<rpcsvc/yp_prot.h>`, either `YP_TRUE` or an error code. See `ypprot_err()`, for a function that converts an NIS name service protocol error code to a `ypclnt` layer error code.

The key and value parameters are somewhat different than defined in the synopsis section above. First, the memory pointed to by the *inkey* and *inval* parameters is private to the `yp_all()` function, and is overwritten with the arrival of each new key-value pair. The `foreach()` function must do something useful with the contents of that memory, but it does not own the memory itself. Key and value objects presented to the `foreach()` function look exactly as they do in the server's map. If they were not NEWLINE-terminated or null-terminated in the map, they would not be here either.

The *indata* parameter is the contents of the `incallback->data` element passed to `yp_all()`. The `data` element of the callback structure can be used to share state information between the `foreach()` function and the mainline code. Its use is optional, and no part of the NIS client package inspects its contents; cast it to something useful, or ignore it. The `foreach()` function is Boolean. It should return 0 to indicate that it wants to be called again for further received key-value pairs, or non-zero to stop the flow of key-value pairs. If `foreach()` returns a non-zero value, it is not called again. The functional value of `yp_all()` is then 0.

`yp_order()`

`yp_order()` returns the order number for a map. This function is not supported if the `ypbind` process on the client's system is bound to an NIS+ server running in "YP-compatibility mode."

`yp_master()`

`yp_master()` returns the machine name of the master NIS server for a map.

`yperr_string()`

`yperr_string()` returns a pointer to an error message string that is null-terminated but contains no period or NEWLINE.

`ypprot_err()`

`ypprot_err()` takes an NIS name service protocol error code as input, and returns a `ypclnt()` layer error code, which can be used as an input to `yperr_string()`.

## RETURN VALUES

All integer functions return 0 if the requested operation is successful, or one of the following errors if the operation fails:

<code>YPERR_ACCESS</code>	Access violation.
<code>YPERR_BADARGS</code>	The arguments to the function are bad.

## ypclnt(3NSL)

YPERR_BADDB	The YP database is bad.
YPERR_BUSY	The database is busy.
YPERR_DOMAIN	Cannot bind to server on this domain.
YPERR_KEY	No such key in map.
YPERR_MAP	No such map in server's domain.
YPERR_NODOM	Local domain name not set.
YPERR_NOMORE	No more records in map database.
YPERR_PMAP	Cannot communicate with rpcbnd.
YPERR_RESRC	Resource allocation failure.
YPERR_RPC	RPC failure; domain has been unbound.
YPERR_YPBIND	Cannot communicate with ypbnd.
YPERR_YPERR	Internal YP server or client error.
YPERR_YPSESV	Cannot communicate with ypserv.
YPERR_VERS	YP version mismatch.

**FILES** /usr/lib/libnsl.so.1

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

**SEE ALSO** nis+(1), ypcat(1), ypmatch(1), ypwhich(1), rpc.nisd(1M), rpcbnd(1M), ypbnd(1M), ypserv(1M), sysinfo(2), malloc(3C), ypfiles(4), attributes(5)

<b>NAME</b>	yp_update – change NIS information				
<b>SYNOPSIS</b>	<pre>#include &lt;rpcsvc/ypclnt.h&gt;  int <b>yp_update</b>(char *domain, char *map, unsigned ypop, char *key, int                keylen, char *data, int datalen);</pre>				
<b>DESCRIPTION</b>	<p>yp_update() is used to make changes to the NIS database. The syntax is the same as that of yp_match() except for the extra parameter <i>ypop</i> which may take on one of four values. If it is POP_CHANGE then the data associated with the key will be changed to the new value. If the key is not found in the database, then yp_update() will return YPERR_KEY. If <i>ypop</i> has the value YPOP_INSERT then the key-value pair will be inserted into the database. The error YPERR_KEY is returned if the key already exists in the database. To store an item into the database without concern for whether it exists already or not, pass <i>ypop</i> as YPOP_STORE and no error will be returned if the key already or does not exist. To delete an entry, the value of <i>ypop</i> should be YPOP_DELETE.</p> <p>This routine depends upon secure RPC, and will not work unless the network is running secure RPC.</p>				
<b>RETURN VALUES</b>	<p>If the value of <i>ypop</i> is POP_CHANGE, yp_update() returns the error YPERR_KEY if the key is not found in the database.</p> <p>If the value of <i>ypop</i> is POP_INSERT, yp_update() returns the error YPERR_KEY if the key already exists in the database.</p>				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Unsafe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Unsafe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Unsafe				
<b>SEE ALSO</b>	secure_rpc(3NSL), ypclnt(3NSL), attributes(5)				
<b>NOTES</b>	This interface is unsafe in multithreaded applications. Unsafe interfaces should be called only from the main thread.				

yp\_update(3NSL)

---

# Index

---

## A

- abandon an LDAP operation in progress —
  - ldap\_abandon, 297
- accept — accept a connection on a socket, 18
- accept a security context initiated by a peer application — gss\_accept\_sec\_context, 202
- acquire a handle for a pre-existing credential by name — gss\_acquire\_cred, 208
- add a credential-element to a credential —
  - gss\_add\_cred, 211
- add an object identifier to an object identifier set
  - gss\_add\_oid\_set\_member, 215
- address in an XFN reference
  - fn\_ref\_addr\_assign, 131
  - fn\_ref\_addr\_copy, 131
  - fn\_ref\_addr\_create, 131
  - fn\_ref\_addr\_data, 131
  - fn\_ref\_addr\_description, 131
  - fn\_ref\_addr\_destroy, 131
  - fn\_ref\_addr\_length, 131
  - FN\_ref\_addr\_t, 131
  - fn\_ref\_addr\_type, 131
- allow application to determine maximum message size with resulting output token of a specified maximum size —
  - gss\_wrap\_size\_limit, 273
- attach a cryptographic message —
  - gss\_wrap, 271
- attribute modifications, list of
  - fn\_attrmodlist\_add, 77
  - fn\_attrmodlist\_assign, 77
  - fn\_attrmodlist\_copy, 77
  - fn\_attrmodlist\_count, 77

- attribute modifications, list of (*continued*)

- fn\_attrmodlist\_create, 77
- fn\_attrmodlist\_destroy, 77
- fn\_attrmodlist\_first, 77
- fn\_attrmodlist\_next, 77
- FN\_attrmodlist\_t, 77

- attribute search options

- fn\_search\_control\_assign, 136
- fn\_search\_control\_copy, 136
- fn\_search\_control\_create, 136
- fn\_search\_control\_destroy, 136
- fn\_search\_control\_follow\_links, 136
- fn\_search\_control\_max\_names, 136
- fn\_search\_control\_return\_attr\_ids, 136
- fn\_search\_control\_return\_ref, 136
- fn\_search\_control\_scope, 136
- FN\_search\_control\_t, 136

- auth\_destroy — library routines for client side remote procedure call authentication, 445

- authnone\_create — library routines for client side remote procedure call authentication, 445

- authsys\_create — library routines for client side remote procedure call authentication, 445
- authsys\_create\_default — library routines for client side remote procedure call authentication, 445

## B

- Basic Encoding Rules library decoding functions

attribute search options (*continued*)

- `ber_alloc_t`, 22
- `ber_bvdup`, 22
- `ber_bvecfree`, 22
- `ber_bvfree`, 22
- `ber_decode`, 22
- `ber_first_element`, 22
- `ber_flatten`, 22
- `ber_free`, 22
- `ber_get_bitstring`, 22
- `ber_get_boolean`, 22
- `ber_get_int`, 22
- `ber_get_next`, 22
- `ber_get_null`, 22
- `ber_get_stringa`, 22
- `ber_get_stringal`, 22
- `ber_get_stringb`, 22
- `ber_init`, 22
- `ber_next_element`, 22
- `ber_peek_tag`, 22
- `ber_scanf`, 22
- `ber_skiptag`, 22

`ber_alloc` — simplified Basic Encoding Rules library encoding functions, 27

`ber_alloc_t` — Basic Encoding Rules library decoding functions, 22

`ber_bvdup` — Basic Encoding Rules library decoding functions, 22

`ber_bvecfree` — Basic Encoding Rules library decoding functions, 22

`ber_bvfree` — Basic Encoding Rules library decoding functions, 22

`ber_decode` — Basic Encoding Rules library decoding functions, 22

`ber_encode` — simplified Basic Encoding Rules library encoding functions, 27

`ber_first_element` — Basic Encoding Rules library decoding functions, 22

`ber_flatten` — Basic Encoding Rules library decoding functions, 22

`ber_free` — Basic Encoding Rules library decoding functions, 22

`ber_get_bitstring` — Basic Encoding Rules library decoding functions, 22

`ber_get_boolean` — Basic Encoding Rules library decoding functions, 22

`ber_get_int` — Basic Encoding Rules library decoding functions, 22

`ber_get_next` — Basic Encoding Rules library decoding functions, 22

`ber_get_null` — Basic Encoding Rules library decoding functions, 22

`ber_get_stringa` — Basic Encoding Rules library decoding functions, 22

`ber_get_stringal` — Basic Encoding Rules library decoding functions, 22

`ber_get_stringb` — Basic Encoding Rules library decoding functions, 22

`ber_init` — Basic Encoding Rules library decoding functions, 22

`ber_next_element` — Basic Encoding Rules library decoding functions, 22

`ber_peek_tag` — Basic Encoding Rules library decoding functions, 22

`ber_printf` — simplified Basic Encoding Rules library encoding functions, 27

`ber_put_bitstring` — simplified Basic Encoding Rules library encoding functions, 27

`ber_put_boolean` — simplified Basic Encoding Rules library encoding functions, 27

`ber_put_int` — simplified Basic Encoding Rules library encoding functions, 27

`ber_put_null` — simplified Basic Encoding Rules library encoding functions, 27

`ber_put_ostring` — simplified Basic Encoding Rules library encoding functions, 27

`ber_put_seq` — simplified Basic Encoding Rules library encoding functions, 27

`ber_put_set` — simplified Basic Encoding Rules library encoding functions, 27

`ber_put_string` — simplified Basic Encoding Rules library encoding functions, 27

`ber_scanf` — Basic Encoding Rules library decoding functions, 22

`ber_skiptag` — Basic Encoding Rules library decoding functions, 22

`ber_start_seq` — simplified Basic Encoding Rules library encoding functions, 27

`ber_start_set` — simplified Basic Encoding Rules library encoding functions, 27

`bind` — bind a name to a socket, 31

bind a reference to a name — `fn_ctx_bind`, 104

bind a reference to a name and associate attributes with named object — `fn_attr_bind`, 60

bind an address to a transport endpoint —  
 t\_bind, 593

byte order, convert values between host and  
 network

- byteorder, 35
- htonl, 35
- htons, 35
- ntohl, 35
- ntohs, 35

## C

change QOP, service for session

- rpc\_gss\_set\_defaults, 474

character string

- fn\_string\_assign, 152
- fn\_string\_bytecount, 152
- fn\_string\_charcount, 152
- fn\_string\_code\_set, 152
- fn\_string\_compare, 152
- fn\_string\_compare\_substring, 152
- fn\_string\_contents, 152
- fn\_string\_copy, 152
- fn\_string\_create, 152
- fn\_string\_destroy, 152
- fn\_string\_from\_contents, 152
- fn\_string\_from\_str, 152
- fn\_string\_from\_str\_n, 152
- fn\_string\_from\_strings, 152
- fn\_string\_from\_substring, 152
- fn\_string\_is\_empty, 152
- fn\_string\_next\_substring, 152
- fn\_string\_prev\_substring, 152
- fn\_string\_str, 152
- FN\_string\_t, 152

cldap\_close — dispose of connectionless LDAP  
 pointer, 36

cldap\_open — LDAP connectionless  
 communication preparation, 37

cldap\_search\_s — connectionless LDAP  
 search, 38

- Retransmission Algorithm, 38

cldap\_setretryinfo — set connectionless LDAP  
 request retransmission parameters, 40

client side remote procedure call authentication,  
 library routines for

- auth\_destroy, 445

client side remote procedure call authentication,  
 library routines for (*continued*)

- authnone\_create, 445
- authsys\_create, 445
- authsys\_create\_default, 445
- rpc\_clnt\_auth, 445

clnt\_call — library routines for client side  
 calls, 447

clnt\_control — library routines for dealing with  
 creation and manipulation of CLIENT  
 handles, 451

clnt\_create — library routines for dealing with  
 creation and manipulation of CLIENT  
 handles, 451

clnt\_create\_timed — library routines for dealing  
 with creation and manipulation of CLIENT  
 handles, 451

clnt\_create\_vers — library routines for dealing  
 with creation and manipulation of CLIENT  
 handles, 451

clnt\_create\_vers\_timed — library routines for  
 dealing with creation and manipulation of  
 CLIENT handles, 451

clnt\_destroy — library routines for dealing with  
 creation and manipulation of CLIENT  
 handles, 451

clnt\_dg\_create — library routines for dealing  
 with creation and manipulation of CLIENT  
 handles, 451

clnt\_door\_create — library routines for dealing  
 with creation and manipulation of CLIENT  
 handles, 451

clnt\_freeres — library routines for client side  
 calls, 447

clnt\_geterr — library routines for client side  
 calls, 447

clnt\_pcreateerror — library routines for dealing  
 with creation and manipulation of CLIENT  
 handles, 451

clnt\_perrno — library routines for client side  
 calls, 447

clnt\_perror — library routines for client side  
 calls, 447

clnt\_raw\_create — library routines for dealing  
 with creation and manipulation of CLIENT  
 handles, 451

clnt\_spcreateerror — library routines for dealing with creation and manipulation of CLIENT handles, 451  
 clnt\_sperrno — library routines for client side calls, 447  
 clnt\_sperror — library routines for client side calls, 447  
 clnt\_tli\_create — library routines for dealing with creation and manipulation of CLIENT handles, 451  
 clnt\_tp\_create — library routines for dealing with creation and manipulation of CLIENT handles, 451  
 clnt\_tp\_create\_timed — library routines for dealing with creation and manipulation of CLIENT handles, 451  
 clnt\_vc\_create — library routines for dealing with creation and manipulation of CLIENT handles, 451  
 close an open SLP handle — SLPclose, 548  
 communications  
   accept a connection on a socket —  
     accept, 18  
   allocate memory for, 590  
   bind a name to a socket — bind, 31  
   create a pair of connected sockets —  
     socketpair, 581  
   create an endpoint for communication —  
     socket, 576  
   get name of peer connected to socket —  
     getpeername, 180  
   get socket name — getsockname, 193  
   initiate a connection on a socket —  
     connect, 41  
   listen for connections on a socket —  
     listen, 365  
   scatter data in order to test the network —  
     spray, 584  
   send a message from a socket — send,  
     sendto, sendmsg, 520  
   shut down part of a full-duplex connection —  
     shutdown, 536  
 compare two internal-form names —  
   gss\_compare\_name, 218  
 component names spanning multiple naming systems  
   — fn\_composite\_name\_append\_comp, 95  
   — fn\_composite\_name\_append\_name, 95  
   component names spanning multiple naming systems (*continued*)  
     — fn\_composite\_name\_assign, 95  
     — fn\_composite\_name\_copy, 95  
     — fn\_composite\_name\_count, 95  
     — fn\_composite\_name\_create, 95  
     — fn\_composite\_name\_delete\_comp, 95  
     — fn\_composite\_name\_destroy, 95  
     — fn\_composite\_name\_first, 95  
     — fn\_composite\_name\_from\_str, 95  
     — fn\_composite\_name\_from\_string, 95  
     — fn\_composite\_name\_insert\_comp, 95  
     — fn\_composite\_name\_insert\_name, 95  
     — fn\_composite\_name\_is\_empty, 95  
     — fn\_composite\_name\_is\_equal, 95  
     — fn\_composite\_name\_is\_prefix, 95  
     — fn\_composite\_name\_is\_suffix, 95  
     — fn\_composite\_name\_last, 95  
     — fn\_composite\_name\_next, 95  
     — fn\_composite\_name\_prefix, 95  
     — fn\_composite\_name\_prepend\_comp, 95  
     — fn\_composite\_name\_prepend\_name, 95  
     — fn\_composite\_name\_prev, 95  
     — fn\_composite\_name\_suffix, 95  
     — FN\_composite\_name\_t, 95  
     — fn\_string\_from\_composite\_name, 95  
 configuration script  
   execute — doconfig, 48  
 connect — initiate a connection on socket, 41  
 connectionless LDAP search —  
   cldap\_search\_s, 38  
 construct a handle to a context object using the given reference —  
   fn\_ctx\_handle\_from\_ref, 116  
 construct equivalent name in same context —  
   fn\_ctx\_equivalent\_name, 108  
 convert a contiguous string name to GSS\_API internal format — gss\_import\_name, 233  
 convert a GSS-API status code to text —  
   gss\_display\_status, 225  
 convert a mechanism name to export form —  
   gss\_export\_name, 228  
 convert a string to an OID —  
   gss\_str\_to\_oid, 264  
 convert an internal name to a mechanism name —  
   gss\_canonicalize\_name, 216  
 convert an OID to a string —  
   gss\_oid\_to\_str, 255



convert internal-form name to text —  
  gss\_display\_name, 223  
create a copy of an internal name —  
  gss\_duplicate\_name, 227  
create a security context using the RPCSEC\_GSS  
  protocol — rpc\_gss\_seccreate, 470  
create an object-identifier set containing no  
  object identifiers —  
  gss\_create\_empty\_oid\_set, 220  
create subcontext and associate attributes —  
  fn\_attr\_create\_subcontext, 61

## D

delete a GSS-API security context —  
  gss\_delete\_sec\_context, 221  
delete attributes — SLPDelAttrs, 549  
deregister the SLP advertisement —  
  SLPDereg, 551  
descriptions of XFN status codes —  
  xfn\_status\_codes, 706  
determine available security mechanisms —  
  gss\_indicate\_mechs, 237  
determine how long a context will remain valid  
  — gss\_context\_time, 219  
dial — establish an outgoing terminal line  
  connection, 46  
discard a credential handle —  
  gss\_release\_cred, 260  
discard an internal-form name —  
  gss\_release\_name, 261  
dispose of connectionless LDAP pointer —  
  cldap\_close, 36  
dn\_comp — resolver routines, 426  
dn\_expand — resolver routines, 426  
doconfig — execute a configuration script, 48

## E

endservent — get service entry, 189  
escapes SLP reserved characters —  
  SLPEscape, 553  
Ethernet address mapping operations  
  — ethers, 58  
ethers — Ethernet address mapping  
  operations, 58

external data representation  
  See XDR, 682

## F

filter expression for attribute search  
  — fn\_search\_filter\_arguments, 139  
  — fn\_search\_filter\_assign, 139  
  — fn\_search\_filter\_copy, 139  
  — fn\_search\_filter\_create, 139  
  — fn\_search\_filter\_destroy, 139  
  — fn\_search\_filter\_expression, 139  
  — FN\_search\_filter\_t, 139  
find service types — SLPFindSrvTypes, 561  
fn\_attr\_bind — bind a reference to a name and  
  associate attributes with named object, 60  
fn\_attr\_create\_subcontext — create subcontext  
  and associate attributes, 61  
fn\_attr\_ext\_search — search for names whose  
  attributes satisfy filter, 62  
fn\_attr\_get — return specified attribute  
  associated with name, 69  
fn\_attr\_get\_ids — get list of attribute  
  identifiers, 70  
fn\_attr\_get\_values — return values of an  
  attribute, 71  
fn\_attr\_modify — modify specified attribute  
  associated with name, 75  
fn\_attr\_multi\_get — return multiple attributes  
  associated with named object, 80  
fn\_attr\_multi\_modify — modify multiple  
  attributes associated with named object, 84  
fn\_attr\_search — search for atomic name with  
  specified attributes in single context, 86  
fn\_attribute\_add — an XFN attribute, 73  
fn\_attribute\_assign — an XFN attribute, 73  
fn\_attribute\_copy — an XFN attribute, 73  
fn\_attribute\_create — an XFN attribute, 73  
fn\_attribute\_destroy — an XFN attribute, 73  
fn\_attribute\_first — an XFN attribute, 73  
fn\_attribute\_identifier — an XFN attribute, 73  
fn\_attribute\_next — an XFN attribute, 73  
fn\_attribute\_remove — an XFN attribute, 73  
fn\_attribute\_syntax — an XFN attribute, 73  
FN\_attribute\_t — an XFN attribute, 73  
fn\_attribute\_valuecount — an XFN  
  attribute, 73

fn\_attrmodlist\_add — a list of attribute modifications, 77  
 fn\_attrmodlist\_assign — a list of attribute modifications, 77  
 fn\_attrmodlist\_copy — a list of attribute modifications, 77  
 fn\_attrmodlist\_count — a list of attribute modifications, 77  
 fn\_attrmodlist\_create — a list of attribute modifications, 77  
 fn\_attrmodlist\_destroy — a list of attribute modifications, 77  
 fn\_attrmodlist\_first — a list of attribute modifications, 77  
 fn\_attrmodlist\_next — a list of attribute modifications, 77  
 FN\_attrmodlist\_t — a list of attribute modifications, 77  
 fn\_attrset\_add — a set of XFN attributes, 91  
 fn\_attrset\_assign — a set of XFN attributes, 91  
 fn\_attrset\_copy — a set of XFN attributes, 91  
 fn\_attrset\_count — a set of XFN attributes, 91  
 fn\_attrset\_create — a set of XFN attributes, 91  
 fn\_attrset\_destroy — a set of XFN attributes, 91  
 fn\_attrset\_first — a set of XFN attributes, 91  
 fn\_attrset\_get — a set of XFN attributes, 91  
 fn\_attrset\_next — a set of XFN attributes, 91  
 fn\_attrset\_remove — a set of XFN attributes, 91  
 FN\_attrset\_t — a set of XFN attributes, 91  
 fn\_bindinglist\_destroy — list the atomic names and references bound in a context, 118  
 fn\_bindinglist\_next — list the atomic names and references bound in a context, 118  
 FN\_bindinglist\_t — list the atomic names and references bound in a context, 118  
 fn\_composite\_name\_append\_comp — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_append\_name — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_assign — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_copy — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_count — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_create — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_delete\_comp — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_destroy — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_first — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_from\_str — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_from\_string — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_insert\_comp — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_insert\_name — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_is\_empty — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_is\_equal — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_is\_prefix — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_is\_suffix — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_last — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_next — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_prefix — component names spanning multiple naming systems, 95  
 fn\_composite\_name\_prepend\_comp — component names spanning multiple naming systems, 95

`fn_composite_name_prepend_name` — component names spanning multiple naming systems, 95  
`fn_composite_name_prev` — component names spanning multiple naming systems, 95  
`fn_composite_name_suffix` — component names spanning multiple naming systems, 95  
`FN_composite_name_t` — component names spanning multiple naming systems, 95  
`fn_compound_name_append_comp` — an XFN compound name, 100  
`fn_compound_name_assign` — an XFN compound name, 100  
`fn_compound_name_copy` — an XFN compound name, 100  
`fn_compound_name_count` — an XFN compound name, 100  
`fn_compound_name_delete_all` — an XFN compound name, 100  
`fn_compound_name_delete_comp` — an XFN compound name, 100  
`fn_compound_name_destroy` — an XFN compound name, 100  
`fn_compound_name_first` — an XFN compound name, 100  
`fn_compound_name_from_syntax_attrs` — an XFN compound name, 100  
`fn_compound_name_get_syntax_attrs` — an XFN compound name, 100  
`fn_compound_name_insert_comp` — an XFN compound name, 100  
`fn_compound_name_is_empty` — an XFN compound name, 100  
`fn_compound_name_is_equal` — an XFN compound name, 100  
`fn_compound_name_is_prefix` — an XFN compound name, 100  
`fn_compound_name_is_suffix` — an XFN compound name, 100  
`fn_compound_name_last` — an XFN compound name, 100  
`fn_compound_name_next` — an XFN compound name, 100  
`fn_compound_name_prefix` — an XFN compound name, 100  
`fn_compound_name_prepend_comp` — an XFN compound name, 100  
`fn_compound_name_prev` — an XFN compound name, 100  
`fn_compound_name_suffix` — an XFN compound name, 100  
`FN_compound_name_t` — an XFN compound name, 100  
`fn_ctx_bind` — bind a reference to a name, 104  
`fn_ctx_equivalent_name` — construct equivalent name in same context, 108  
`fn_ctx_handle_from_initial` — return a handle to the Initial Context, 114  
`fn_ctx_handle_from_ref` — construct a handle to a context object using the given reference, 116  
`fn_ctx_list_bindings` — list the atomic names and references bound in a context, 118  
`fn_ctx_list_names` — list the atomic names bound in a context, 119  
`fn_ctx_lookup_link` — look up the link reference bound to a name, 123  
`fn_ctx_rename` — rename the name of a binding, 124  
`FN_ctx_t` — an XFN context, 127  
`fn_ext_searchlist_destroy` — search for names whose attributes satisfy filter, 62  
`fn_ext_searchlist_next` — search for names whose attributes satisfy filter, 62  
`FN_ext_searchlist_t` — search for names whose attributes satisfy filter, 62  
`FN_identifier_t` — an XFN identifier, 130  
`fn_multigetlist_destroy` — return multiple attributes associated with named object, 80  
`fn_multigetlist_next` — return multiple attributes associated with named object, 80  
`FN_multigetlist_t` — return multiple attributes associated with named object, 80  
`fn_namelist_destroy` — list the atomic names bound in a context, 119  
`fn_namelist_next` — list the atomic names bound in a context, 119  
`FN_namelist_t` — list the atomic names bound in a context, 119  
`fn_ref_addr_assign` — an address in an XFN reference, 131  
`fn_ref_addr_copy` — an address in an XFN reference, 131  
`fn_ref_addr_create` — an address in an XFN reference, 131

fn\_ref\_addr\_data — an address in an XFN reference, 131  
 fn\_ref\_addr\_description — an address in an XFN reference, 131  
 fn\_ref\_addr\_destroy — an address in an XFN reference, 131  
 fn\_ref\_addr\_length — an address in an XFN reference, 131  
 FN\_ref\_addr\_t — an address in an XFN reference, 131  
 fn\_ref\_addr\_type — an address in an XFN reference, 131  
 fn\_ref\_addrcount — an XFN reference, 133  
 fn\_ref\_append\_addr — an XFN reference, 133  
 fn\_ref\_assign — an XFN reference, 133  
 fn\_ref\_copy — an XFN reference, 133  
 fn\_ref\_create — an XFN reference, 133  
 fn\_ref\_create\_link — an XFN reference, 133  
 fn\_ref\_delete\_addr — an XFN reference, 133  
 fn\_ref\_delete\_all — an XFN reference, 133  
 fn\_ref\_description — an XFN reference, 133  
 fn\_ref\_destroy — an XFN reference, 133  
 fn\_ref\_first — an XFN reference, 133  
 fn\_ref\_insert\_addr — an XFN reference, 133  
 fn\_ref\_is\_link — an XFN reference, 133  
 fn\_ref\_link\_name — an XFN reference, 133  
 fn\_ref\_next — an XFN reference, 133  
 fn\_ref\_prepend\_addr — an XFN reference, 133  
 FN\_ref\_t — an XFN reference, 133  
 fn\_ref\_type — an XFN reference, 133  
 fn\_search\_control\_assign — options for attribute search, 136  
 fn\_search\_control\_copy — options for attribute search, 136  
 fn\_search\_control\_create — options for attribute search, 136  
 fn\_search\_control\_destroy — options for attribute search, 136  
 fn\_search\_control\_follow\_links — options for attribute search, 136  
 fn\_search\_control\_max\_names — options for attribute search, 136  
 fn\_search\_control\_return\_attr\_ids — options for attribute search, 136  
 fn\_search\_control\_return\_ref — options for attribute search, 136  
 fn\_search\_control\_scope — options for attribute search, 136  
 FN\_search\_control\_t — options for attribute search, 136  
 fn\_search\_filter\_arguments — filter expression for attribute search, 139  
 fn\_search\_filter\_assign — filter expression for attribute search, 139  
 fn\_search\_filter\_copy — filter expression for attribute search, 139  
 fn\_search\_filter\_create — filter expression for attribute search, 139  
 fn\_search\_filter\_destroy — filter expression for attribute search, 139  
 fn\_search\_filter\_expression — filter expression for attribute search, 139  
 FN\_search\_filter\_t — filter expression for attribute search, 139  
     BNF of Filter Expression, 140  
     Extended Operations, 142  
     Precedence, 140  
     Relational Operators, 141  
     Specification of Filter Expression, 140  
     Wildcarded Strings, 141  
 fn\_searchlist\_destroy — terminate search for atomic name with specified attributes in single context, 86  
 fn\_searchlist\_next — search for next atomic name with specified attributes in single context, 86  
 FN\_searchlist\_t — search for atomic name with specified attributes in single context, 86  
 fn\_status\_advance\_by\_name — an XFN status object, 147  
 fn\_status\_append\_remaining\_name — an XFN status object, 147  
 fn\_status\_append\_resolved\_name — an XFN status object, 147  
 fn\_status\_assign — an XFN status object, 147  
 fn\_status\_code — an XFN status object, 147  
 fn\_status\_copy — an XFN status object, 147  
 fn\_status\_create — an XFN status object, 147  
 fn\_status\_description — an XFN status object, 147  
 fn\_status\_destroy — an XFN status object, 147  
 fn\_status\_diagnostic\_message — an XFN status object, 147  
 fn\_status\_is\_success — an XFN status object, 147

fn\_status\_link\_code — an XFN status object, 147  
 fn\_status\_link\_diagnostic\_message — an XFN status object, 147  
 fn\_status\_link\_remaining\_name — an XFN status object, 147  
 fn\_status\_link\_resolved\_name — an XFN status object, 147  
 fn\_status\_link\_resolved\_ref — an XFN status object, 147  
 fn\_status\_remaining\_name — an XFN status object, 147  
 fn\_status\_resolved\_name — an XFN status object, 147  
 fn\_status\_resolved\_ref — an XFN status object, 147  
 fn\_status\_set — an XFN status object, 147  
 fn\_status\_set\_code — an XFN status object, 147  
 fn\_status\_set\_diagnostic\_message — an XFN status object, 147  
 fn\_status\_set\_link\_code — an XFN status object, 147  
 fn\_status\_set\_link\_diagnostic\_message — an XFN status object, 147  
 fn\_status\_set\_link\_remaining\_name — an XFN status object, 147  
 fn\_status\_set\_link\_resolved\_name — an XFN status object, 147  
 fn\_status\_set\_link\_resolved\_ref — an XFN status object, 147  
 fn\_status\_set\_remaining\_name — an XFN status object, 147  
 fn\_status\_set\_resolved\_name — an XFN status object, 147  
 fn\_status\_set\_resolved\_ref — an XFN status object, 147  
 fn\_status\_set\_success — an XFN status object, 147  
 FN\_status\_t — an XFN status object, 147  
 fn\_string\_assign — a character string, 152  
 fn\_string\_bytecount — a character string, 152  
 fn\_string\_charcount — a character string, 152  
 fn\_string\_code\_set — a character string, 152  
 fn\_string\_compare — a character string, 152  
 fn\_string\_compare\_substring — a character string, 152  
 fn\_string\_contents — a character string, 152  
 fn\_string\_copy — a character string, 152  
 fn\_string\_create — a character string, 152  
 fn\_string\_destroy — a character string, 152  
 fn\_string\_from\_composite\_name — component names spanning multiple naming systems, 95  
 fn\_string\_from\_compound\_name — an XFN compound name, 100  
 fn\_string\_from\_contents — a character string, 152  
 fn\_string\_from\_str — a character string, 152  
 fn\_string\_from\_str\_n — a character string, 152  
 fn\_string\_from\_strings — a character string, 152  
 fn\_string\_from\_substring — a character string, 152  
 fn\_string\_is\_empty — a character string, 152  
 fn\_string\_next\_substring — a character string, 152  
 fn\_string\_prev\_substring — a character string, 152  
 fn\_string\_str — a character string, 152  
 FN\_string\_t — a character string, 152  
 fn\_valuelist\_destroy — return values of an attribute, 71  
 fn\_valuelist\_next — return values of an attribute, 71  
 FN\_valuelist\_t — return values of an attribute, 71  
 FNS  
 component names spanning multiple naming systems  
   *See* FN\_composite\_name\_t  
 fn\_attr\_bind — bind a reference to a name and associate attributes with named object, 60  
 fn\_attr\_create\_subcontext — create subcontext and associate attributes, 61  
 fn\_attr\_ext\_search — search for names whose attributes satisfy filter, 62  
 fn\_attr\_search — search for atomic name with specified attributes in single context, 86  
 fn\_ctx\_equivalent\_name — construct equivalent name in same context, 108  
 fn\_ext\_searchlist\_destroy — search for names whose attributes satisfy filter, 62

FNS (*continued*)

fn\_ext\_searchlist\_next — search for names whose attributes satisfy filter, 62  
FN\_ext\_searchlist\_t — search for names whose attributes satisfy filter, 62  
FN\_search\_control\_t — options for attribute search, 136  
FN\_search\_filter\_t — filter expression for attribute search, 139  
fn\_searchlist\_destroy — terminate search for atomic name with specified attributes in single context, 86  
fn\_searchlist\_next — search for next atomic name with specified attributes in single context, 86  
FN\_searchlist\_t — search for atomic name with specified attributes in single context, 86  
fp\_resstat — resolver routines, 426  
free a BerElement structure from memory — ldap\_ber\_free, 300  
free buffer storage allocated by a GSS-API function — gss\_release\_buffer, 259  
free memory allocated by LDAP API functions — ldap\_memfree, 343  
free storage associated with a GSS-API-generated gss\_OID\_set object — gss\_release\_oid\_set, 263  
freeaddrinfo — translate between node name and address, 155  
freehostent — get IP node entry, 166  
frees memory — SLPFree, 563  
functions to map Internet Protocol network interface names and interface indexes — if\_freenameindex, 278  
functions to map Internet Protocol network interface names and interface indexes — if\_indextoname, 278  
functions to map Internet Protocol network interface names and interface indexes — if\_nameindex, 278  
functions to map Internet Protocol network interface names and interface indexes — if\_nametoindex, 278

**G**

gai\_strerror — translate between node name and address, 155  
generic transport name-to-address translation — netdir, 368  
— netdir\_free, 368  
— netdir\_getbyaddr, 368  
— netdir\_getbyname, 368  
— netdir\_mergeaddr, 368  
— netdir\_options, 368  
— netdir\_perror, 368  
— netdir\_sperror, 368  
— taddr2uaddr, 368  
— uaddr2taddr, 368  
get IP node entry — freehostent, 166  
get IP node entry — getipnodebyaddr, 166  
get IP node entry — getipnodebyname, 166  
get service entry — getservbyname, 189  
endservent, 189  
getservbyname\_r, 189  
getservbyport, 189  
getservbyport\_r, 189  
getservent, 189  
getservent\_r, 189  
setservent, 189  
get credentials of client — rpc\_gss\_getcred, 460  
get error codes on failure — rpc\_gss\_get\_error, 462  
get list of attribute identifiers — fn\_attr\_get\_ids, 70  
get maximum data length for transmission — rpc\_gss\_max\_data\_length, 467  
— rpc\_gss\_svc\_max\_data\_length, 467  
get principal names at server — rpc\_get\_principal\_name, 465  
getaddrinfo — translate between node name and address, 155  
getipnodebyaddr — get IP node entry, 166  
getipnodebyname — get IP node entry, 166  
getnameinfo — translate between node name and address, 155  
getpeername — get name of peer connected to socket, 180  
getpublickey — retrieve public or secret key, 185  
getsecretkey — retrieve public or secret key, 185

getservbyname — get service entry, 189  
 getservbyname\_r — get service entry, 189  
 getservbyport — get service entry, 189  
 getservbyport\_r — get service entry, 189  
 getservent — get service entry, 189  
 getservent\_r — get service entry, 189  
 gss\_accept\_sec\_context — accept a security context initiated by a peer application, 202  
 gss\_acquire\_cred — acquire a handle for a pre-existing credential by name, 208  
 gss\_add\_cred — add a credential-element to a credential, 211  
 gss\_add\_oid\_set\_member — add an object identifier to an object identifier set, 215  
 gss\_canonicalize\_name — convert an internal name to a mechanism name, 216  
 gss\_compare\_name — compare two internal-form names, 218  
 gss\_context\_time — determine how long a context will remain valid, 219  
 gss\_create\_empty\_oid\_set — create an object-identifier set containing no object identifiers, 220  
 gss\_delete\_sec\_context — delete a GSS-API security context, 221  
 gss\_display\_name — convert internal-form name to text, 223  
 gss\_display\_status — convert a GSS-API status code to text, 225  
 gss\_duplicate\_name — create a copy of an internal name, 227  
 gss\_export\_name — convert a mechanism name to export form, 228  
 gss\_export\_sec\_context — transfer a security context to another process, 229  
 gss\_import\_name — convert a contiguous string name to GSS\_API internal format, 233  
 gss\_import\_sec\_context — import security context established by another process, 235  
 gss\_indicate\_mechs — determine available security mechanisms, 237  
 gss\_init\_sec\_context — initiate a GSS-API security context with a peer application, 238  
 gss\_inquire\_context — obtain information about a security context, 245

gss\_inquire\_cred — obtain information about a credential, 248  
 gss\_inquire\_cred\_by\_mech — obtain per-mechanism information about a credential, 250  
 gss\_inquire\_mechs\_for\_name — list mechanisms that support the specified name-type, 252  
 gss\_inquire\_names\_for\_mech — list the name-types supported by the specified mechanism, 254  
 gss\_oid\_to\_str — convert an OID to a string, 255  
 gss\_process\_context\_token — pass asynchronous token to security service, 257  
 gss\_release\_buffer — free buffer storage allocated by a GSS-API function, 259  
 gss\_release\_cred — discard a credential handle, 260  
 gss\_release\_name — discard an internal-form name, 261  
 gss\_release\_oid — release an object identifier, 262  
 gss\_release\_oid\_set — free storage associated with a GSS-API-generated gss\_OID\_set object, 263  
 gss\_str\_to\_oid — convert a string to an OID, 264  
 gss\_test\_oid\_set\_member — interrogate an object identifier set, 266  
 gss\_verify\_mic — verify integrity of a received message, 269  
 gss\_wrap — attach a cryptographic message, 271  
 gss\_wrap — verify a message with attached cryptographic message, 267  
 gss\_wrap\_size\_limit — allow application to determine maximum message size with resulting output token of a specified maximum size, 273

## H

host machines, remote  
     return information about users — rusers, rusers, 514  
 hostalias — resolver routines, 426

hstrerror — resolver routines, 426

## I

if\_freenameindex — functions to map Internet Protocol network interface names and interface indexes, 278  
if\_indextoname — functions to map Internet Protocol network interface names and interface indexes, 278  
if\_nameindex — functions to map Internet Protocol network interface names and interface indexes, 278  
if\_nametoindex — functions to map Internet Protocol network interface names and interface indexes, 278  
import security context established by another process — gss\_import\_sec\_context, 235  
inet — Internet address manipulation, 280  
inet\_addr — Internet address manipulation, 280  
inet\_lnaof — Internet address manipulation, 280  
inet\_makeaddr — Internet address manipulation, 280  
inet\_netof — Internet address manipulation, 280  
inet\_network — Internet address manipulation, 280  
inet\_ntoa — Internet address manipulation, 280  
inet\_ntop — Internet address manipulation, 280  
inet\_pton — Internet address manipulation, 280  
inet6 — Internet address manipulation, 280  
initialize an LDAP session — ldap\_init, 348  
initialize an LDAP session — ldap\_open, 348  
initiate a GSS-API security context with a peer application — gss\_init\_sec\_context, 238  
Internet address manipulation — inet6, 280  
Internet address manipulation — inet\_addr, 280  
Internet address manipulation — inet\_lnaof, 280  
Internet address manipulation — inet\_makeaddr, 280

Internet address manipulation — inet\_netof, 280  
Internet address manipulation — inet\_network, 280  
Internet address manipulation — inet\_ntoa, 280  
Internet address manipulation — inet\_ntop, 280  
Internet address manipulation — inet, 280  
Internet address manipulation — inet\_pton, 280  
interrogate an object identifier set — gss\_test\_oid\_set\_member, 266

## L

LDAP bind functions — ldap\_bind, 301  
LDAP bind functions — ldap\_bind\_s, 301  
LDAP bind functions — ldap\_sasl\_bind, 301  
LDAP bind functions — ldap\_sasl\_bind\_s, 301  
LDAP bind functions — ldap\_set\_rebind\_proc, 301  
LDAP bind functions — ldap\_simple\_bind, 301  
LDAP bind functions — ldap\_simple\_bind\_s, 301  
LDAP bind functions — ldap\_unbind\_ext, 301  
LDAP bind functions — ldap\_unbind, 301  
LDAP bind functions — ldap\_unbind\_s, 301  
LDAP client caching functions — ldap\_memcache\_destroy, 340  
LDAP client caching functions — ldap\_memcache\_flush, 340  
LDAP client caching functions — ldap\_memcache\_get, 340  
LDAP client caching functions — ldap\_memcache\_init, 340  
LDAP client caching functions — ldap\_memcache, 340  
LDAP client caching functions — ldap\_memcache\_set, 340  
LDAP client caching functions — ldap\_memcache\_update, 340  
LDAP entry parsing and counting functions — ldap\_count\_entries, 324  
LDAP entry parsing and counting functions — ldap\_count\_references, 324



LDAP entry parsing and counting functions —  
  ldap\_first\_entry, 324

LDAP entry parsing and counting functions —  
  ldap\_first\_reference, 324

LDAP entry parsing and counting functions —  
  ldap\_next\_entry, 324

LDAP entry parsing and counting functions —  
  ldap\_next\_reference, 324

LDAP filter generating functions —  
  ldap\_build\_filter, 331

LDAP filter generating functions —  
  ldap\_getfilter\_free, 331

LDAP filter generating functions —  
  ldap\_getfilter, 331

LDAP filter generating functions —  
  ldap\_getfirstfilter, 331

LDAP filter generating functions —  
  ldap\_getnextfilter, 331

LDAP filter generating functions —  
  ldap\_init\_getfilter\_buf, 331

LDAP filter generating functions —  
  ldap\_init\_getfilter, 331

LDAP filter generating functions —  
  ldap\_setfilteraffixes, 331

LDAP protocol error handling functions —  
  ldap\_err2string, 319

LDAP protocol error handling functions —  
  ldap\_error, 319

LDAP protocol error handling functions —  
  ldap\_perror, 319

LDAP protocol error handling functions —  
  ldap\_result2error, 319

ldap — Lightweight Directory Access Protocol  
  package  
  Index, 288

ldap\_8859\_to\_t61 — LDAP character set  
  translation functions, 304

ldap\_abandon — abandon an LDAP operation  
  in progress, 297

ldap\_add — perform an LDAP add  
  operation, 298

ldap\_add\_ext — perform an LDAP add  
  operation, 298

ldap\_add\_ext\_s — perform an LDAP add  
  operation, 298

ldap\_add\_s — perform an LDAP add  
  operation, 298

LDAP attribute remapping functions  
  — ldap\_free\_friendlymap, 327  
  — ldap\_friendly\_name, 327

LDAP attribute value handling functions  
  — ldap\_count\_values, 338  
  — ldap\_get\_values, 338  
  — ldap\_get\_values\_len, 338

ldap\_ber\_free — free a BerElement structure  
  from memory, 300

ldap\_bind — LDAP bind functions, 301  
  General Authentication, 301  
  Rebinding While Following Referral, 302  
  Simple Authentication, 301  
  Unbinding, 302

ldap\_bind\_s — LDAP bind functions, 301

ldap\_build\_filter — LDAP filter generating  
  functions, 331

LDAP character set translation functions  
  — ldap\_8859\_to\_t61, 304  
  — ldap\_enable\_translation, 304  
  — ldap\_set\_string\_translators, 304  
  — ldap\_t61\_to\_8859, 304  
  — ldap\_translate\_from\_t61, 304  
  — ldap\_translate\_to\_t61, 304

ldap\_compare — LDAP compare  
  operation, 306

ldap\_compare\_ext — LDAP compare  
  operation, 306

ldap\_compare\_ext\_s — LDAP compare  
  operation, 306

LDAP compare operation  
  — ldap\_compare, 306  
  — ldap\_compare\_ext, 306  
  — ldap\_compare\_ext\_s, 306  
  — ldap\_compare\_s, 306

ldap\_compare\_s — LDAP compare  
  operation, 306

LDAP connectionless communication  
  preparation — cldap\_open, 37

LDAP control disposal  
  — ldap\_control\_free, 308  
  — ldap\_controls\_free, 308

ldap\_control\_free — LDAP control  
  disposal, 308

ldap\_controls\_free — LDAP control  
  disposal, 308

ldap\_count\_entries — LDAP entry parsing and  
  counting functions, 324

ldap\_count\_messages — LDAP message processing functions, 326  
 ldap\_count\_references — LDAP entry parsing and counting functions, 324  
 ldap\_count\_values — LDAP attribute value handling functions, 338  
 ldap\_delete — LDAP delete operation, 309  
 ldap\_delete\_ext — LDAP delete operation, 309  
 ldap\_delete\_ext\_s — LDAP delete operation, 309  
 LDAP delete operation  
   — ldap\_delete, 309  
   — ldap\_delete\_ext, 309  
   — ldap\_delete\_ext\_s, 309  
   — ldap\_delete\_s, 309  
 ldap\_delete\_s — LDAP delete operation, 309  
 LDAP display template functions  
   — ldap\_disptmpl, 310  
   — ldap\_first\_disptmpl, 310  
   — ldap\_first\_tmplcol, 310  
   — ldap\_first\_tmplrow, 310  
   — ldap\_free\_templates, 310  
   — ldap\_init\_templates, 310  
   — ldap\_init\_templates\_buf, 310  
   — ldap\_next\_disptmpl, 310  
   — ldap\_next\_tmplcol, 310  
   — ldap\_next\_tmplrow, 310  
   — ldap\_oc2template, 310  
   — ldap\_tmplattrs, 310  
 ldap\_disptmpl — LDAP display template functions, 310  
   DISPTMPL Structure Elements, 312  
   Syntax IDs, 314  
   TMPLITEM Structure Elements, 313  
 LDAP DN handling functions  
   — ldap\_dn2ufn, 328  
   — ldap\_dns\_to\_dn, 328  
   — ldap\_explode\_dn, 328  
   — ldap\_explode\_dns, 328  
   — ldap\_get\_dn, 328  
   — ldap\_is\_dns\_dn, 328  
 ldap\_dn\_to\_url — LDAP Uniform Resource Locator functions, 361  
 ldap\_dn2ufn — LDAP DN handling functions, 328  
 ldap\_dns\_to\_dn — LDAP DN handling functions, 328  
 ldap\_dns\_to\_url — LDAP Uniform Resource Locator functions, 361  
 ldap\_enable\_translation — LDAP character set translation functions, 304  
 LDAP entry display functions  
   — ldap\_entry2text, 316  
   — ldap\_entry2text\_search, 316  
   — ldap\_vals2text, 316  
 LDAP entry modification functions  
   — ldap\_modify, 344  
   — ldap\_modify\_ext, 344  
   — ldap\_modify\_ext\_s, 344  
   — ldap\_modify\_s, 344  
 LDAP entry sorting functions  
   — ldap\_sort, 357  
   — ldap\_sort\_entries, 357  
   — ldap\_sort\_strcasecmp, 357  
   — ldap\_sort\_values, 357  
 ldap\_entry2text — LDAP entry display functions, 316  
 ldap\_entry2text\_search — LDAP entry display functions, 316  
 ldap\_err2string — LDAP protocol error handling functions, 319  
 ldap\_error — LDAP protocol error handling functions, 319  
 ldap\_explode\_dn — LDAP DN handling functions, 328  
 ldap\_explode\_dns — LDAP DN handling functions, 328  
 ldap\_first\_attribute — step through LDAP entry attributes, 323  
 ldap\_first\_disptmpl — LDAP display template functions, 310  
 ldap\_first\_entry — LDAP entry parsing and counting functions, 324  
 ldap\_first\_message — LDAP message processing functions, 326  
 ldap\_first\_reference — LDAP entry parsing and counting functions, 324  
 ldap\_first\_searchobj — LDAP search preference configuration routines, 355  
 ldap\_first\_tmplcol — LDAP display template functions, 310  
 ldap\_first\_tmplrow — LDAP display template functions, 310  
 ldap\_free\_friendlymap — LDAP attribute remapping functions, 327

ldap\_free\_searchprefs — LDAP search preference configuration routines, 355  
 ldap\_free\_templates — LDAP display template functions, 310  
 ldap\_free\_urldesc — LDAP Uniform Resource Locator functions, 361  
 ldap\_friendly\_name — LDAP attribute remapping functions, 327  
 ldap\_get\_dn — LDAP DN handling functions, 328  
 ldap\_get\_lang\_values — return an attribute's values that matches a specified language subtype, 333  
 ldap\_get\_lang\_values\_len — return an attribute's values that matches a specified language subtype, 333  
 ldap\_get\_values — LDAP attribute value handling functions, 338  
 ldap\_get\_values\_len — LDAP attribute value handling functions, 338  
 ldap\_getfilter — LDAP filter generating functions, 331  
 ldap\_getfilter\_free — LDAP filter generating functions, 331  
 ldap\_getfirstfilter — LDAP filter generating functions, 331  
 ldap\_getnextfilter — LDAP filter generating functions, 331  
 ldap\_init — initialize an LDAP session, 348  
 ldap\_init\_getfilter — LDAP filter generating functions, 331  
 ldap\_init\_getfilter\_buf — LDAP filter generating functions, 331  
 ldap\_init\_searchprefs — LDAP search preference configuration routines, 355  
 ldap\_init\_searchprefs\_buf — LDAP search preference configuration routines, 355  
 ldap\_init\_templates — LDAP display template functions, 310  
 ldap\_init\_templates\_buf — LDAP display template functions, 310  
 ldap\_is\_dns\_dn — LDAP DN handling functions, 328  
 ldap\_is\_ldap\_url — LDAP Uniform Resource Locator functions, 361  
 ldap\_memcache — LDAP client caching functions, 340  
 ldap\_memcache\_destroy — LDAP client caching functions, 340  
 ldap\_memcache\_flush — LDAP client caching functions, 340  
 ldap\_memcache\_get — LDAP client caching functions, 340  
 ldap\_memcache\_init — LDAP client caching functions, 340  
 ldap\_memcache\_set — LDAP client caching functions, 340  
 ldap\_memcache\_update — LDAP client caching functions, 340  
 ldap\_memfree — free memory allocated by LDAP API functions, 343  
 LDAP message processing functions  
   — ldap\_count\_message, 326  
   — ldap\_first\_message, 326  
   — ldap\_msgtype, 326  
   — ldap\_next\_message, 326  
 LDAP message result parser  
   — ldap\_parse\_extended\_result, 350  
   — ldap\_parse\_result, 350  
   — ldap\_parse\_sasl\_bind\_result, 350  
 ldap\_modify — LDAP entry modification functions, 335, 344  
 ldap\_modify\_ext — LDAP entry modification functions, 344  
 ldap\_modify\_ext\_s — LDAP entry modification functions, 344  
 ldap\_modify\_s — LDAP entry modification functions, 344  
 ldap\_modrdn — modify LDAP entry RDN, 346  
 ldap\_modrdn\_s — modify LDAP entry RDN, 346  
 ldap\_modrdn2 — modify LDAP entry RDN, 346  
 ldap\_modrdn2\_s — modify LDAP entry RDN, 346  
 ldap\_msgfree — wait for and return LDAP operation result, 351  
 ldap\_msgtype — LDAP message processing functions, 326  
 ldap\_next\_attribute — step through LDAP entry attributes, 323  
 ldap\_next\_disptmpl — LDAP display template functions, 310

ldap\_next\_entry — LDAP entry parsing and counting functions, 324  
 ldap\_next\_message — LDAP message processing functions, 326  
 ldap\_next\_reference — LDAP entry parsing and counting functions, 324  
 ldap\_next\_searchobj — LDAP search preference configuration routines, 355  
 ldap\_next\_tmplcol — LDAP display template functions, 310  
 ldap\_next\_tmplrow — LDAP display template functions, 310  
 ldap\_oc2template — LDAP display template functions, 310  
 ldap\_open — initialize an LDAP session, 348  
 ldap\_parse\_extended\_result — LDAP message result parser, 350  
 ldap\_parse\_result — LDAP message result parser, 350  
 ldap\_parse\_sasl\_bind\_result — LDAP message result parser, 350  
 ldap\_perror — LDAP protocol error handling functions, 319  
 ldap\_rename — modify LDAP entry RDN, 346  
 ldap\_rename\_s — modify LDAP entry RDN, 346  
 ldap\_result — wait for and return LDAP operation result, 351  
 ldap\_result2error — LDAP protocol error handling functions, 319  
 ldap\_sasl\_bind — LDAP bind functions, 301  
 ldap\_sasl\_bind\_s — LDAP bind functions, 301  
 ldap\_search — LDAP search operations, 353  
 ldap\_search\_ext — LDAP search operations, 353  
 ldap\_search\_ext\_s — LDAP search operations, 353  
 LDAP search operations  
   — ldap\_search, 353  
   — ldap\_search\_ext, 353  
   — ldap\_search\_ext\_s, 353  
   — ldap\_search\_s, 353  
   — ldap\_search\_st, 353  
 LDAP search preference configuration routines  
   — ldap\_first\_searchobj, 355  
   — ldap\_free\_searchprefs, 355  
   — ldap\_init\_searchprefs, 355  
   — ldap\_init\_searchprefs\_buf, 355  
   — ldap\_next\_searchobj, 355  
   — ldap\_searchprefs, 355  
 ldap\_search\_s — LDAP search operations, 353  
 ldap\_search\_st — LDAP search operations, 353  
 ldap\_searchprefs — LDAP search preference configuration routines, 355  
 ldap\_set\_rebind\_proc — LDAP bind functions, 301  
 ldap\_set\_string\_translators — LDAP character set translation functions, 304  
 ldap\_setfilteraffixes — LDAP filter generating functions, 331  
 ldap\_simple\_bind — LDAP bind functions, 301  
 ldap\_simple\_bind\_s — LDAP bind functions, 301  
 ldap\_sort — LDAP entry sorting functions, 357  
 ldap\_sort\_entries — LDAP entry sorting functions, 357  
 ldap\_sort\_strcasecmp — LDAP entry sorting functions, 357  
 ldap\_sort\_values — LDAP entry sorting functions, 357  
 ldap\_t61\_to\_8859 — LDAP character set translation functions, 304  
 ldap\_tmplattrs — LDAP display template functions, 310  
 ldap\_translate\_from\_t61 — LDAP character set translation functions, 304  
 ldap\_translate\_to\_t61 — LDAP character set translation functions, 304  
 ldap\_ufn — LDAP user friendly search functions, 359  
 ldap\_ufn\_search\_c — LDAP user friendly search functions, 359  
 ldap\_ufn\_search\_ct — LDAP user friendly search functions, 359  
 ldap\_ufn\_search\_s — LDAP user friendly search functions, 359  
 ldap\_ufn\_setfilter — LDAP user friendly search functions, 359  
 ldap\_ufn\_setprefix — LDAP user friendly search functions, 359  
 LDAP search preference configuration routines (*continued*)

ldap\_ufn\_timeout — LDAP user friendly search functions, 359  
 ldap\_unbind — LDAP bind functions, 301  
 ldap\_unbind\_ext — LDAP bind functions, 301  
 ldap\_unbind\_s — LDAP bind functions, 301  
 LDAP Uniform Resource Locator functions  
   — ldap\_dn\_to\_url, 361  
   — ldap\_dns\_to\_url, 361  
   — ldap\_free\_urldesc, 361  
   — ldap\_is\_ldap\_url, 361  
   — ldap\_url, 361  
   — ldap\_url\_parse, 361  
   — ldap\_url\_search, 361  
   — ldap\_url\_search\_s, 361  
   — ldap\_url\_search\_st, 361  
 ldap\_url — LDAP Uniform Resource Locator functions, 361  
 ldap\_url\_parse — LDAP Uniform Resource Locator functions, 361  
 ldap\_url\_search — LDAP Uniform Resource Locator functions, 361  
 ldap\_url\_search\_s — LDAP Uniform Resource Locator functions, 361  
 ldap\_url\_search\_st — LDAP Uniform Resource Locator functions, 361  
 LDAP user friendly search functions  
   — ldap\_ufn, 359  
   — ldap\_ufn\_search\_c, 359  
   — ldap\_ufn\_search\_ct, 359  
   — ldap\_ufn\_search\_s, 359  
   — ldap\_ufn\_setfilter, 359  
   — ldap\_ufn\_setprefix, 359  
   — ldap\_ufn\_timeout, 359  
 ldap\_vals2text — LDAP entry display functions, 316  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_control, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_create, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_create\_timed, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_create\_vers, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_create\_vers\_timed, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_dg\_create, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_door\_create, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_pcreateerror, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_raw\_create, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_spcreateerror, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_tli\_create, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_tp\_create, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_tp\_create\_timed, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — clnt\_vc\_create, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — rpc\_clnt\_create, 451  
 library routines for dealing with creation and manipulation of CLIENT handles — rpc\_createerr, 451  
 library routines for the creation of server handles — rpc\_svc\_create, 500  
 library routines for the creation of server handles — svc\_control, 500  
 library routines for the creation of server handles — svc\_create, 500  
 library routines for the creation of server handles — svc\_destroy, 500

- library routines for the creation of server handles — `svc_dg_create`, 500
- library routines for the creation of server handles — `svc_door_create`, 500
- library routines for the creation of server handles — `svc_fd_create`, 500
- library routines for the creation of server handles — `svc_raw_create`, 500
- library routines for the creation of server handles — `svc_tli_create`, 500
- library routines for the creation of server handles — `svc_tp_create`, 500
- library routines for the creation of server handles — `svc_vc_create`, 500
- library routines for client side calls
  - `clnt_call`, 447
  - `clnt_freeres`, 447
  - `clnt_geterr`, 447
  - `clnt_perrno`, 447
  - `clnt_perror`, 447
  - `clnt_sperrno`, 447
  - `clnt_sperror`, 447
  - `rpc_broadcast`, 447
  - `rpc_broadcast_exp`, 447
  - `rpc_call`, 447
  - `rpc_clnt_calls`, 447
- library routines for RPC servers
  - `rpc_svc_calls`, 496
  - `svc_dg_enablecache`, 496
  - `svc_done`, 496
  - `svc_exit`, 496
  - `svc_fdset`, 496
  - `svc_freeargs`, 496
  - `svc_getargs`, 496
  - `svc_getreq_common`, 496
  - `svc_getreq_poll`, 496
  - `svc_getreqset`, 496
  - `svc_getrpcaller`, 496
  - `svc_max_pollfd`, 496
  - `svc_pollfd`, 496
  - `svc_run`, 496
  - `svc_sendreply`, 496
- list mechanisms that support the specified name-type —
  - `gss_inquire_mechs_for_name`, 252
- list the atomic names and references bound in a context
  - `fn_bindinglist_destroy`, 118

- list the atomic names and references bound in a context (*continued*)
  - `fn_bindinglist_next`, 118
  - `FN_bindinglist_t`, 118
  - `fn_ctx_list_bindings`, 118
- list the atomic names bound in a context
  - `fn_ctx_list_names`, 119
  - `fn_namelist_destroy`, 119
  - `fn_namelist_next`, 119
  - `FN_namelist_t`, 119
- list the name-types supported by the specified mechanism —
  - `gss_inquire_names_for_mech`, 254
- listen — listen for connections on a socket, 365
- look up the link reference bound to a name —
  - `fn_ctx_lookup_link`, 123

## M

- map ASCII mechanism to OID
  - `rpc_gss_mech_to_oid`, 463, 468
- map ASCII qop to number
  - `rpc_gss_qop_to_num`, 463, 468
- map SLP error codes to messages —
  - `slp_strerror`, 573
- modify LDAP entry RDN
  - `ldap_modrdn`, 346
  - `ldap_modrdn_s`, 346
  - `ldap_modrdn2`, 346
  - `ldap_modrdn2_s`, 346
  - `ldap_rename`, 346
  - `ldap_rename_s`, 346
- modify multiple attributes associated with named object — `fn_attr_multi_modify`, 84
- modify specified attribute associated with name
  - `fn_attr_modify`, 75

## N

- `netdir` — generic transport name-to-address translation, 368
- `netdir_free` — generic transport name-to-address translation, 368
- `netdir_getbyaddr` — generic transport name-to-address translation, 368

netdir\_getbyname — generic transport  
   name-to-address translation, 368  
 netdir\_mergeaddr — generic transport  
   name-to-address translation, 368  
 netdir\_options — generic transport  
   name-to-address translation, 368  
 netdir\_perror — generic transport  
   name-to-address translation, 368  
 netdir\_spperror — generic transport  
   name-to-address translation, 368  
 network configuration database entry  
   — endnetconfig, 176  
   — freenetconfigent, 176  
   — getnetconfig, 176  
   — getnetconfigent, 176  
   — nc\_perror, 176  
   — nc\_spperror, 176  
   — setnetconfig, 176  
 network configuration entry corresponding to  
   NETPATH  
   — endnetpath, 178  
   — getnetpath, 178  
   — setnetpath, 178  
 network entry  
   — endnetent, 172  
   — getnetbyaddr, 172  
   — getnetbyaddr\_r, 172  
   — getnetbyname, 172  
   — getnetbyname\_r, 172  
   — getnetent, 172  
   — getnetent\_r, 172  
   — setnetent, 172  
 network host entry  
   — endhostent, 159  
   — gethostbyaddr, 159  
   — gethostbyaddr\_r, 159  
   — gethostbyname, 159  
   — gethostbyname\_r, 159  
   — gethostent, 159  
   — gethostent\_r, 159  
   — sethostent, 159  
 network listener service  
   format and send listener service request  
     message — nlsrequest, 410  
   get client's data passed via the listener —  
     nlsgetcall, 408  
   get name of transport provider —  
     nlsprovider, 409  
 network protocol entry  
   — endprotoent, 182  
   — getprotobyname, 182  
   — getprotobyname\_r, 182  
   — getprotobynumber, 182  
   — getprotobynumber\_r, 182  
   — getprotoent, 182  
   — getprotoent\_r, 182  
   — setprotoent, 182  
 NIS+ table functions — nis\_tables  
   nis\_first\_entry, 399  
   nis\_modify\_entry, 399  
   nis\_next\_entry, 399  
   nis\_remove\_entry, 399  
 NIS, change information  
   — yp\_update, 715  
 NIS+ error messages  
   nis\_error, 372  
   nis\_lerror, 372  
   nis\_perror, 372  
   nis\_spperror, 372  
   nis\_spperror\_r, 372  
 NIS+ group manipulation functions  
   — nis\_addmember, 374  
   — nis\_creategroup, 374  
   — nis\_destroygroup, 374  
   — nis\_groups, 374  
   — nis\_ismember, 374  
   — nis\_print\_group\_entry, 374  
   — nis\_removemember, 374  
   — nis\_verifygroup, 374  
 NIS+ local names  
   — nis\_freenames, 396  
   — nis\_getnames, 396  
   — nis\_local\_directory, 377  
   — nis\_local\_group, 377  
   — nis\_local\_host, 377  
   — nis\_local\_names, 377  
   — nis\_local\_principal, 377  
 NIS+ log administration functions  
   — nis\_checkpoint, 393  
   — nis\_ping, 393  
 NIS+ miscellaneous functions  
   — nis\_freesservelist, 394  
   — nis\_freetags, 394  
   — nis\_getservlist, 394  
   — nis\_mkdir, 394

NIS+ miscellaneous functions (*continued*)

- nis\_rmdir, 394
- nis\_server, 394
- nis\_servstate, 394
- nis\_stats, 394

NIS+ namespace functions

- nis\_add, 379
- nis\_freeresult, 379
- nis\_lookup, 379
- nis\_modify, 379
- nis\_names, 379
- nis\_remove, 379

NIS+ object formats

- nis\_objects, 385

NIS+ subroutines

- nis\_clone\_object, 396
- nis\_destroy\_object, 396
- nis\_dir\_cmp, 396
- nis\_domain\_of, 396
- nis\_leaf\_of, 396
- nis\_name\_of, 396
- nis\_print\_object, 396
- nis\_subr, 396

NIS+ table functions

- nis\_add\_entry, 399
- nis\_first\_entry, 399
- nis\_list, 399
- nis\_modify\_entry, 399
- nis\_next\_entry, 399
- nis\_remove\_entry, 399
- nis\_tables, 399

NIS client interface

- yp\_all, 710
- yp\_bind, 710
- yp\_first, 710
- yp\_get\_default\_domain, 710
- yp\_master, 710
- yp\_match, 710
- yp\_next, 710
- yp\_order, 710
- yp\_unbind, 710
- ypcnt, 710
- yperr\_string, 710
- ypprot\_err, 710

nis\_tables — NIS+ table functions, 399

## O

- obtain information about a credential —
  - gss\_inquire\_cred, 248
- obtain information about a security context —
  - gss\_inquire\_context, 245
- obtain per-mechanism information about a credential —
  - gss\_inquire\_cred\_by\_mech, 250
- open an SLP handle — SLPOpen, 566
- overview of the XFN interface — xfn, 695
- an overview of XFN attribute operations —
  - xfn\_attributes, 696
- XFN compound syntax: an overview of XFN model for compound name parsing —
  - xfn\_compound\_names, 700

## P

- parse service URL — SLPParseSrvURL, 568
- pass asynchronous token to security service —
  - gss\_process\_context\_token, 257
- perform an LDAP add operation
  - ldap\_add, 298
  - ldap\_add\_ext, 298
  - ldap\_add\_ext\_s, 298
  - ldap\_add\_s, 298
- publickey — retrieve public or secret key, 185

## R

- rac\_drop() — remote asynchronous calls, 477
- rac\_poll() — remote asynchronous calls, 477
- rac\_recv() — remote asynchronous calls, 477
- rac\_send() — remote asynchronous calls, 477
- rcmd — routines for returning a stream to a remote command, 412
- rcmd\_af — routines for returning a stream to a remote command, 412
- receive a message from a socket — recv, 414
  - recvfrom, 414
  - recvmsg, 414
- rcv — receive a message from a socket, 414
- rcvfrom — receive a message from a socket, 414
- recvmsg — receive a message from a socket, 414



- register an SLP advertisement — SLPReg, 570
- release an object identifier —
  - gss\_release\_oid, 262
- remote procedure calls, library routines for —
  - rpc, 434
- remote system
  - return information about users — rusers,
    - rnusers, 514
  - write to — rstat, 513
  - write to — rwall, 515
- rename the name of a binding —
  - fn\_ctx\_rename, 124
- res\_hostalias — resolver routines, 426
- res\_init — resolver routines, 426
- res\_mkquery — resolver routines, 426
- res\_nclose — resolver routines, 426
- res\_ninit — resolver routines, 426
- res\_nmquery — resolver routines, 426
- res\_nquery — resolver routines, 426
- res\_nquerydomain — resolver routines, 426
- res\_nsearch — resolver routines, 426
- res\_nsend — resolver routines, 426
- res\_nsendsigned — resolver routines, 426
- res\_query — resolver routines, 426
- res\_querydomain — resolver routines, 426
- res\_search — resolver routines, 426
- res\_send — resolver routines, 426
- res\_update — resolver routines, 426
- resolver — resolver routines, 426
- resolver routines — dn\_comp, 426
- resolver routines — dn\_expand, 426
- resolver routines — fp\_resstat, 426
- resolver routines — hostalias, 426
- resolver routines — hstrerror, 426
- resolver routines — res\_hostalias, 426
- resolver routines — res\_init, 426
- resolver routines — res\_mkquery, 426
- resolver routines — res\_nclose, 426
- resolver routines — res\_ninit, 426
- resolver routines — res\_nmquery, 426
- resolver routines — res\_nquerydomain, 426
- resolver routines — res\_nquery, 426
- resolver routines — res\_nsearch, 426
- resolver routines — res\_nsend, 426
- resolver routines — res\_nsendsigned, 426
- resolver routines — resolver, 426
- resolver routines — res\_querydomain, 426
- resolver routines — res\_query, 426
- resolver routines — res\_search, 426
- resolver routines — res\_send, 426
- resolver routines — res\_update, 426
- retrieve public or secret key —
  - getpublickey, 185
  - getsecretkey, 185
  - publickey, 185
- return an attribute's values that matches a specified language subtype —
  - ldap\_get\_lang\_values, 333
- return an attribute's values that matches a specified language subtype —
  - ldap\_get\_lang\_values\_len, 333
- return stream to a remote command —
  - rexec\_af, 432
- return stream to a remote command —
  - rexec, 432
- return a handle to the Initial Context —
  - fn\_ctx\_handle\_from\_initial, 114
- return list of configured and discovered scopes —
  - SLPFindScopes, 557
- return multiple attributes associated with named object
  - fn\_attr\_multi\_get, 80
  - fn\_multigetlist\_destroy, 80
  - fn\_multigetlist\_next, 80
  - FN\_multigetlist\_t, 80
- return service attributes — SLPFindAttrs, 555
- return service URLs — SLPFindSrvs, 559
- return SLP configuration property —
  - SLPGetProperty, 564
- return specified attribute associated with name —
  - fn\_attr\_get, 69
- return the maximum allowed refresh interval —
  - SLPGetRefreshInterval, 565
- return values of an attribute
  - fn\_attr\_get\_values, 71
  - fn\_valuelist\_destroy, 71
  - fn\_valuelist\_next, 71
  - FN\_valuelist\_t, 71
- rexec — return stream to a remote command, 432
- rexec\_af — return stream to a remote command, 432
- rnusers — return information about users on remote machines, 514
- routines for returning a stream to a remote command — rcmd\_af, 412

routines for returning a stream to a remote command — `rcmd`, 412  
 routines for returning a stream to a remote command — `rresvport_af`, 412  
 routines for returning a stream to a remote command — `rresvport`, 412  
 routines for returning a stream to a remote command — `ruserok`, 412  
`rpc` — library routines for remote procedure calls, 434  
 RPC  
   data transmission using XDR routines — `xdr`, 682  
 RPC, XDR library routines  
   — `rpc_xdr`, 511  
   — `xdr_accepted_reply`, 511  
   — `xdr_authsys_parms`, 511  
   — `xdr_callhdr`, 511  
   — `xdr_callmsg`, 511  
   — `xdr_opaque_auth`, 511  
   — `xdr_rejected_reply`, 511  
   — `xdr_replymsg`, 511  
 RPC bind service library routines  
   — `rpc_getmaps`, 443  
   — `rpcb_getaddr`, 443  
   — `rpcb_gettime`, 443  
   — `rpcb_rmtcall`, 443  
   — `rpcb_set`, 443  
   — `rpcb_unset`, 443  
   — `rpcbind`, 443  
`rpc_broadcast` — library routines for client side calls, 447  
`rpc_broadcast_exp` — library routines for client side calls, 447  
`rpc_call` — library routines for client side calls, 447  
`rpc_clnt_auth` — library routines for client side remote procedure call authentication, 445  
`rpc_clnt_calls` — library routines for client side calls, 447  
   Routines, 447  
`rpc_clnt_create` — library routines for dealing with creation and manipulation of CLIENT handles, 451  
   Routines, 452  
`rpc_createerr` — library routines for dealing with creation and manipulation of CLIENT handles, 451

RPC entry  
   — `endrpcent`, 186  
   — `getrpcbyname`, 186  
   — `getrpcbyname_r`, 186  
   — `getrpcbynumber`, 186  
   — `getrpcbynumber_r`, 186  
   — `getrpcent`, 186  
   — `getrpcent_r`, 186  
   — `setrpcent`, 186  
`rpc_gss_getcred` — get credentials of client, 460  
`rpc_gss_seccreate` — create a security context using the RPCSEC\_GSS protocol, 470  
 RPC library routine for manipulating global RPC attributes for client and server applications  
   — `rpc_control`, 458  
 RPC library routines for registering servers  
   — `rpc_reg`, 509  
   — `rpc_svc_reg`, 509  
   — `svc_auth_reg`, 509  
   — `svc_reg`, 509  
   — `svc_unreg`, 509  
   — `xprt_register`, 509  
   — `xprt_unregister`, 509  
 RPC library routines for server side errors  
   — `rpc_svc_err`, 505  
   — `svcerr_auth`, 505  
   — `svcerr_decode`, 505  
   — `svcerr_noproc`, 505  
   — `svcerr_noprog`, 505  
   — `svcerr_progvers`, 505  
   — `svcerr_systemerr`, 505  
   — `svcerr_weakauth`, 505  
 RPC obsolete library routines  
   — `authdes_create`, 487  
   — `authunix_create_default`, 487  
   — `callrpc`, 487  
   — `clnt_broadcast`, 487  
   — `clntraw_create`, 487  
   — `clnttcp_create`, 487  
   — `clntudp_bufcreate`, 487  
   — `clntudp_create`, 487  
   — `get_myaddress`, 487  
   — `getrpcport`, 487  
   — `pmap_getmaps`, 487  
   — `pmap_getport`, 487  
   — `pmap_rmtcall`, 487

RPC obsolete library routines (*continued*)

- pmap\_set, 487
  - pmap\_unset, 487
  - registerrpc, 487
  - rpc\_soc, 487
  - svc\_fds, 487
  - svc\_getcaller, 487
  - svc\_getreq, 487
  - svc\_register, 487
  - svc\_unregister, 487
  - svcd\_create, 487
  - svcraw\_create, 487
  - svctcp\_create, 487
  - svcudp\_bufcreate, 487
  - svcudp\_create, 487
  - xdr\_authunix\_parms, 487
- rpc routines
- rac\_drop() — remote asynchronous calls, 477
  - rac\_poll() — remote asynchronous calls, 477
  - rac\_recv() — remote asynchronous calls, 477
  - rac\_send() — remote asynchronous calls, 477
- rpc\_svc\_calls — library routines for RPC servers, 496
- Routines, 496
- rpc\_svc\_create — library routines for the creation of server handles, 500
- rpc — security flavor incorporating GSS-API onto ONC RPC, 481
- rresvport — routines for returning a stream to a remote command, 412
- rresvport\_af — routines for returning a stream to a remote command, 412
- rstat — get performance data from remote kernel, 513
- ruserok — routines for returning a stream to a remote command, 412
- users — return information about users on remote machines, 514
- xdr\_utmpidlearr, 514
- rwall — write to specified remote machines, 515

**S**

- search for atomic name with specified attributes in single context
- fn\_attr\_search, 86
  - fn\_searchlist\_destroy, 86
  - fn\_searchlist\_next, 86
  - FN\_searchlist\_t, 86
- search for names whose attributes satisfy filter
- fn\_attr\_ext\_search, 62
  - fn\_ext\_searchlist\_destroy, 62
  - fn\_ext\_searchlist\_next, 62
  - FN\_ext\_searchlist\_t, 62
- send — send message from a socket, 520
- sendmsg — send message from a socket, 520
- sendto — send message from a socket, 520
- Service Access Facility library function
- doconfig, 48
- Service Location Protocol Application Programming Interface — slp\_api, 538
- set an SLP configuration property — SLPSetProperty, 572
- set connectionless LDAP request retransmission parameters — ldap\_setretryinfo, 40
- set server principal name
- rpc\_gss\_set\_svc\_name, 475
- setservent — get service entry, 189
- shutdown — shut down part of a full-duplex connection, 536
- simplified Basic Encoding Rules library
- encoding functions
- ber\_alloc, 27
  - ber\_encode, 27
  - ber\_printf, 27
  - ber\_put\_bitstring, 27
  - ber\_put\_boolean, 27
  - ber\_put\_int, 27
  - ber\_put\_null, 27
  - ber\_put\_ostring, 27
  - ber\_put\_seq, 27
  - ber\_put\_set, 27
  - ber\_put\_string, 27
  - ber\_start\_seq, 27
  - ber\_start\_set, 27
- slp\_api — Service Location Protocol Application Programming Interface, 538
- slp\_strerror — map SLP error codes to messages, 573
- SLPclose — close an open SLP handle, 548

SLPDelAttrs — delete attributes, 549

SLPDereg — deregister the SLP advertisement, 551

SLPEscape — escapes SLP reserved characters, 553

SLPFindAttrs — return service attributes, 555

SLPFindScopes — return list of configured and discovered scopes, 557

SLPFindSrvs — return service URLs, 559

SLPFindSrvTypes — find service types, 561

SLPFree — frees memory, 563

SLPGetProperty — return SLP configuration property, 564

SLPGetRefreshInterval — return the maximum allowed refresh interval, 565

SLPOpen — open an SLP handle, 566

SLPParseSrvURL — parse service URL, 568

SLPReg — register an SLP advertisement, 570

SLPSetProperty — set an SLP configuration property, 572

SLPUnescape — translate escaped characters into UTF-8, 574

socket — create an endpoint for communication, 576

socket

- accept a connection — accept, 18
- bind a name — bind, 31
- get options — getsockopt, 195
- get name — getsockname, 193
- get name of connected peer — getpeername, 180
- initiate a connection — connect, 41
- listen for connections — listen, 365
- send message from — send, sendto, sendmsg, 520
- set options — setsockopt, 195
- shut down part of a full-duplex connection — shutdown, 536

socketpair — create a pair of connected sockets, 581

spray — scatter data in order to test the network, 584

step through LDAP entry attributes

- ldap\_first\_attribute, 323
- ldap\_next\_attribute, 323

STREAMS

- accept a connection on a socket — accept, 18

STREAMS (*continued*)

- bind a name to a socket — bind, 31
- create a pair of connected sockets — socketpair, 581
- create an endpoint for communication — socket, 576
- get and set socket options — getsockopt, setsockopt, 195
- get name of peer connected to socket — getpeername, 180
- get socket name — getsockname, 193
- initiate a connection on a socket — connect, 41
- listen for connections on a socket — listen, 365
- send a message from a socket — send, sendto, sendmsg, 520
- shut down part of a full-duplex connection — shutdown, 536
- svc\_control — library routines for the creation of server handles, 500
- svc\_create — library routines for the creation of server handles, 500
- svc\_destroy — library routines for the creation of server handles, 500
- svc\_dg\_create — library routines for the creation of server handles, 500
- svc\_dg\_enablecache — library routines for RPC servers, 496
- svc\_done — library routines for RPC servers, 496
- svc\_door\_create — library routines for the creation of server handles, 500
- svc\_exit — library routines for RPC servers, 496
- svc\_fd\_create — library routines for the creation of server handles, 500
- svc\_fdset — library routines for RPC servers, 496
- svc\_freeargs — library routines for RPC servers, 496
- svc\_getargs — library routines for RPC servers, 496
- svc\_getreq\_common — library routines for RPC servers, 496
- svc\_getreq\_poll — library routines for RPC servers, 496

svc\_getreqset — library routines for RPC servers, 496  
 svc\_getrpccaller — library routines for RPC servers, 496  
 svc\_max\_pollfd — library routines for RPC servers, 496  
 svc\_pollfd — library routines for RPC servers, 496  
 svc\_raw\_create — library routines for the creation of server handles, 500  
 svc\_run — library routines for RPC servers, 496  
 svc\_sendreply — library routines for RPC servers, 496  
 svc\_tli\_create — library routines for the creation of server handles, 500  
 svc\_tp\_create — library routines for the creation of server handles, 500  
 svc\_vc\_create — library routines for the creation of server handles, 500

## T

t\_alloc — allocate memory for argument structures, 590  
 t\_bind — bind an address to a transport endpoint, 593  
 taddr2uaddr — generic transport name-to-address translation, 368  
 terminal line  
   establish an outgoing connection — dial, 46  
 transfer a security context to another process — gss\_export\_sec\_context, 229  
 translate between node name and address — freeaddrinfo, 155  
 translate between node name and address — gai\_strerror, 155  
 translate between node name and address — getaddrinfo, 155  
 translate between node name and address — getnameinfo, 155  
 translate escaped characters into UTF-8 — SLPUnescape, 574  
 transport functions  
   allocate memory, 590

## U

uaddr2taddr — generic transport name-to-address translation, 368  
 users  
   return information from remote machines — rusers, rnusers, 514

## V

verify a message with attached cryptographic message — gss\_wrap, 267  
 verify integrity of a received message — gss\_verify\_mic, 269

## W

wait for and return LDAP operation result — ldap\_msgfree, 351  
 wait for and return LDAP operation result — ldap\_result, 351

## X

XDR library routines  
   — xdr, 682  
   — xdr\_admin, 684  
   — xdr\_control, 684  
   — xdr\_getpos, 684  
   — xdr\_inline, 684  
   — xdr\_setpos, 684  
   — xdr\_sizeof, 684  
   — xdrrec\_endofrecord, 684  
   — xdrrec\_eof, 684  
   — xdrrec\_readbytes, 684  
   — xdrrec\_skiprecord, 684  
 XDR library routines for complex data structures  
   — xdr\_array, 686  
   — xdr\_bytes, 686  
   — xdr\_complex, 686  
   — xdr\_opaque, 686  
   — xdr\_pointer, 686  
   — xdr\_reference, 686  
   — xdr\_string, 686  
   — xdr\_union, 686

XDR library routines for complex data structures (*continued*)

- xdr\_vector, 686
- xdr\_wrapstring, 686

XDR library routines for RPC

- rpc\_xdr, 511
- xdr\_accepted\_reply, 511
- xdr\_authsys\_parms, 511
- xdr\_callhdr, 511
- xdr\_callmsg, 511
- xdr\_opaque\_auth, 511
- xdr\_rejected\_reply, 511
- xdr\_replymsg, 511

XDR library routines for simple data structures

- xdr\_bool, 691
- xdr\_char, 691
- xdr\_double, 691
- xdr\_enum, 691
- xdr\_float, 691
- xdr\_free, 691
- xdr\_hyper, 691
- xdr\_int, 691
- xdr\_long, 691
- xdr\_longlong\_t, 691
- xdr\_quadruple, 691
- xdr\_short, 691
- xdr\_simple, 691
- xdr\_u\_char, 691
- xdr\_u\_hyper, 691
- xdr\_u\_int, 691
- xdr\_u\_long, 691
- xdr\_u\_longlong\_t, 691
- xdr\_u\_short, 691
- xdr\_void, 691

xdr\_statstime — get performance data from remote kernel, 513

xdr\_statsvar — get performance data from remote kernel, 513

XDR stream creation library routines

- xdr\_create, 689
- xdr\_destroy, 689
- xdrmem\_create, 689
- xdrrec\_create, 689
- xdrstdio\_create, 689

xfn — overview of the XFN interface, 695

XFN attribute

- fn\_attribute\_add, 73
- fn\_attribute\_assign, 73

XFN attribute (*continued*)

- fn\_attribute\_copy, 73
- fn\_attribute\_create, 73
- fn\_attribute\_destroy, 73
- fn\_attribute\_first, 73
- fn\_attribute\_identifier, 73
- fn\_attribute\_next, 73
- fn\_attribute\_remove, 73
- fn\_attribute\_syntax, 73
- FN\_attribute\_t, 73
- fn\_attribute\_valuecount, 73

xfn\_attributes — an overview of XFN attribute operations, 696

XFN attributes, a set of

- fn\_attrset\_add, 91
- fn\_attrset\_assign, 91
- fn\_attrset\_copy, 91
- fn\_attrset\_count, 91
- fn\_attrset\_create, 91
- fn\_attrset\_destroy, 91
- fn\_attrset\_first, 91
- fn\_attrset\_get, 91
- fn\_attrset\_next, 91
- fn\_attrset\_remove, 91
- FN\_attrset\_t, 91

XFN compound name

- fn\_compound\_name\_append\_comp, 100
- fn\_compound\_name\_assign, 100
- fn\_compound\_name\_copy, 100
- fn\_compound\_name\_count, 100
- fn\_compound\_name\_delete\_all, 100
- fn\_compound\_name\_delete\_comp, 100
- fn\_compound\_name\_destroy, 100
- fn\_compound\_name\_first, 100
- 
- fn\_compound\_name\_from\_syntax\_attrs, 100
- 
- fn\_compound\_name\_get\_syntax\_attrs, 100
- fn\_compound\_name\_insert\_comp, 100
- fn\_compound\_name\_is\_empty, 100
- fn\_compound\_name\_is\_equal, 100
- fn\_compound\_name\_is\_prefix, 100
- fn\_compound\_name\_is\_suffix, 100
- fn\_compound\_name\_last, 100
- fn\_compound\_name\_next, 100
- fn\_compound\_name\_prefix, 100
- 
- fn\_compound\_name\_prepend\_comp, 100

XFN compound name (*continued*)

- `fn_compound_name_prev`, 100
- `fn_compound_name_suffix`, 100
- `FN_compound_name_t`, 100
- `fn_string_from_compound_name`, 100

`xfn_compound_names` — XFN compound syntax: an overview of XFN model for compound name parsing, 700

an XFN context — `FN_ctx_t`, 127

an XFN identifier — `FN_identifier_t`, 130

XFN reference

- `fn_ref_addrcount`, 133
- `fn_ref_append_addr`, 133
- `fn_ref_assign`, 133
- `fn_ref_copy`, 133
- `fn_ref_create`, 133
- `fn_ref_create_link`, 133
- `fn_ref_delete_addr`, 133
- `fn_ref_delete_all`, 133
- `fn_ref_description`, 133
- `fn_ref_destroy`, 133
- `fn_ref_first`, 133
- `fn_ref_insert_addr`, 133
- `fn_ref_is_link`, 133
- `fn_ref_link_name`, 133
- `fn_ref_next`, 133
- `fn_ref_prepend_addr`, 133
- `FN_ref_t`, 133
- `fn_ref_type`, 133

`xfn_status_codes` — descriptions of XFN status codes, 706

XFN Status Codes, 706

XFN status object

- `fn_status_advance_by_name`, 147
- `fn_status_append_remaining_name`, 147
- `fn_status_append_resolved_name`, 147
- `fn_status_assign`, 147
- `fn_status_code`, 147
- `fn_status_copy`, 147
- `fn_status_create`, 147
- `fn_status_description`, 147
- `fn_status_destroy`, 147
- `fn_status_diagnostic_message`, 147
- `fn_status_is_success`, 147
- `fn_status_link_code`, 147
- `fn_status_link_diagnostic_message`, 147
- `fn_status_link_remaining_name`, 147
- `fn_status_link_resolved_name`, 147

XFN status object (*continued*)

- `fn_status_link_resolved_ref`, 147
- `fn_status_remaining_name`, 147
- `fn_status_resolved_name`, 147
- `fn_status_resolved_ref`, 147
- `fn_status_set`, 147
- `fn_status_set_code`, 147
- `fn_status_set_diagnostic_message`, 147
- `fn_status_set_link_code`, 147
- `fn_status_set_link_diagnostic_message`, 147
- `fn_status_set_link_remaining_name`, 147
- `fn_status_set_link_resolved_name`, 147
- `fn_status_set_link_resolved_ref`, 147
- `fn_status_set_remaining_name`, 147
- `fn_status_set_resolved_name`, 147
- `fn_status_set_resolved_ref`, 147
- `fn_status_set_success`, 147
- `FN_status_t`, 147

