

Programming Interfaces Guide

Copyright © 2001, 2012, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Preface	11
1 Memory and CPU Management	15
Memory Management Interfaces	15
Creating and Using Mappings	15
Removing Mappings	16
Cache Control	16
Library-Level Dynamic Memory	17
Dynamic Memory Allocation	17
Dynamic Memory Debugging	18
Other Memory Control Interfaces	19
CPU Performance Counters	20
API Additions to libcpcc	20
2 Session Description Protocol API	25
Session Description API Overview	25
SDP Library Functions	28
Creating the SDP Session Structure	28
Searching the SDP Session Structure	35
Shutting Down the SDP Session Structure	37
SDP API Utility Functions	38
3 Process Scheduler	43
Overview of the Scheduler	43
Time-Sharing Class	45
System Class	46
Real-time Class	46

Interactive Class	46
Fair-Share Class	46
Fixed-Priority Class	47
Commands and Interfaces	47
priocntl Usage	49
priocntl Interface	50
Interactions With Other Interfaces	50
Kernel Processes	50
Using fork and exec	50
Using nice	51
init(1M)	51
Scheduling and System Performance	51
Process State Transition	51
4 Locality Group APIs	55
Locality Groups Overview	55
Verifying the Interface Version	58
Initializing the Locality Group Interface	58
Using lgrp_init()	59
Using lgrp_fini()	59
Locality Group Hierarchy	60
Using lgrp_cookie_stale()	60
Using lgrp_view()	60
Using lgrp_nlgrps()	61
Using lgrp_root()	61
Using lgrp_parents()	61
Using lgrp_children()	62
Locality Group Contents	62
Using lgrp_resources()	62
Using lgrp_cpus()	63
Using lgrp_mem_size()	63
Locality Group Characteristics	64
Using lgrp_latency_cookie()	64
Locality Groups and Thread and Memory Placement	65
Using lgrp_home()	65

Using <code>madvise()</code>	66
Using <code>madv.so.1</code>	67
Using <code>meminfo()</code>	69
Locality Group Affinity	71
Examples of API Usage	73
5 Input/Output Interfaces	81
Files and I/O Interfaces	81
Basic File I/O	82
Advanced File I/O	83
File System Control	84
Using File and Record Locking	84
Choosing a Lock Type	84
Selecting Advisory or Mandatory Locking	85
Cautions About Mandatory Locking	86
Supported File Systems	86
Terminal I/O Functions	90
6 Interprocess Communication	93
Pipes Between Processes	93
Named Pipes	95
Sockets Overview	95
Doors Overview	95
POSIX Interprocess Communication	96
POSIX Messages	96
POSIX Semaphores	97
POSIX Shared Memory	97
System V IPC	98
Permissions for Messages, Semaphores, and Shared Memory	98
IPC Interfaces, Key Arguments, and Creation Flags	98
System V Messages	99
System V Semaphores	101
System V Shared Memory	105

7 Socket Interfaces	109
Overview of Sockets	109
Socket Libraries	110
Socket Types	110
Interface Sets	111
Socket Basics	112
Socket Creation	112
Binding Local Names	113
Connection Establishment	114
Connection Errors	115
Data Transfer	115
Closing Sockets	116
Connecting Stream Sockets	116
Input/Output Multiplexing	121
Datagram Sockets	123
Standard Routines	127
Host and Service Names	127
Host Names – hostent	129
Network Names – netent	129
Protocol Names – protoent	130
Service Names – servent	130
Other Routines	130
Client-Server Programs	131
Sockets and Servers	131
Sockets and Clients	133
Connectionless Servers	134
Advanced Socket Topics	136
Out-of-Band Data	136
Nonblocking Sockets	137
Asynchronous Socket I/O	138
Interrupt-Driven Socket I/O	139
Signals and Process Group ID	139
Selecting Specific Protocols	140
Address Binding	140
Socket Options	142
inetd Daemon	143

Broadcasting and Determining Network Configuration	144
Using Multicast	147
Sending IPv4 Multicast Datagrams	148
Receiving IPv4 Multicast Datagrams	149
Sending IPv6 Multicast Datagrams	150
Receiving IPv6 Multicast Datagrams	152
Stream Control Transmission Protocol	153
SCTP Stack Implementation	153
SCTP Socket Interfaces	154
Code Examples of SCTP Use	178
8 Programming With XTI and TLI	189
What Are XTI and TLI?	189
XTI/TLI Read/Write Interface	191
Write Data	192
Read Data	192
Close Connection	192
Advanced XTI/TLI Topics	193
Asynchronous Execution Mode	193
Advanced XTI/TLI Programming Example	194
Asynchronous Networking	199
Networking Programming Models	199
Asynchronous Connectionless-Mode Service	200
Asynchronous Connection-Mode Service	201
Asynchronous Open	202
State Transitions	204
XTI/TLI States	204
Outgoing Events	204
Incoming Events	205
State Tables	206
Guidelines to Protocol Independence	209
XTI/TLI Versus Socket Interfaces	210
Socket-to-XTI/TLI Equivalents	211
Additions to the XTI Interface	213

9 Packet Filtering Hooks	215
Packet Filtering Hooks Interfaces	215
Packet Filtering Hooks Kernel Functions	215
Packet Filtering Hooks Data Types	217
Using the Packet Filtering Hooks Interfaces	218
IP Instances	218
Protocol Registration	219
Event Registration	220
The Packet Hook	222
Packet Filtering Hooks Example	223
10 Transport Selection and Name-to-Address Mapping	243
Transport Selection	243
Name-to-Address Mapping	244
st_raddr .so Library	245
Using the Name-to-Address Mapping Routines	246
11 Real-time Programming and Administration	251
Basic Rules of Real-time Applications	251
Factors that Degrade Response Time	252
Runaway Real-time Processes	254
Asynchronous I/O Behavior	254
The Real-Time Scheduler	255
Dispatch Latency	255
Interface Calls That Control Scheduling	261
Utilities That Control Scheduling	262
Configuring Scheduling	263
Memory Locking	265
Locking a Page	265
Unlocking a Page	266
Locking All Pages	266
Recovering Sticky Locks	266
High Performance I/O	266
POSIX Asynchronous I/O	267
Oracle Solaris Asynchronous I/O	268

Synchronized I/O	270
Interprocess Communication	271
Processing Signals	271
Pipes, Named Pipes, and Message Queues	271
Using Semaphores	272
Shared Memory	272
Asynchronous Network Communication	272
Modes of Networking	272
Timing Facilities	273
Timestamp Interfaces	273
Interval Timer Interfaces	273
12 The Oracle Solaris ABI and ABI Tools	277
What is the Oracle Solaris ABI?	277
Defining the Oracle Solaris ABI	278
Symbol Versioning in Oracle Solaris Libraries	278
Using Symbol Versioning to Label the Oracle Solaris ABI	279
Oracle Solaris ABI Tools	280
appcert Utility	280
What appcert Checks	280
What appcert Does Not Check	281
Working with appcert	282
Using appcert for Application Triage	283
appcert Results	284
Using appt race for Application Verification	286
A UNIX Domain Sockets	289
Creating Sockets	289
Local Name Binding	290
Establishing a Connection	290
Index	293

Preface

The *Programming Interfaces Guide* describes the Oracle Solaris 11 network and system interfaces used by application developers.

SunOS 5.11 is the core of the Oracle Solaris 11 Operating System (Oracle Solaris OS), and conforms to the third edition of the System V Interface Description (SVID) and to the Single UNIX Specification, version 3 (SUSv3). SunOS 5.11 is fully compatible with UNIX System V, Release 4 (SVR4), and supports all System V network services.

Note – This Oracle Solaris release supports systems that use the SPARC and x86 families of processor architectures. The supported systems appear in the Oracle Solaris OS: Hardware Compatibility Lists <http://www.oracle.com/webfolder/technetwork/hcl/index.html>. This document cites any implementation differences between the platform types.

Audience

This book is intended for programmers who are new to the Oracle Solaris platform or want more familiarity with some portion of the interfaces provided. Additional interfaces and facilities for networked applications are described in the *ONC+ Developer's Guide*.

This manual assumes a familiarity with secure programming techniques, as communication with other systems or processes provides avenues for hackers to launch attacks. [Appendix A](#), “Secure Coding Guidelines for Developers,” in *Developer's Guide to Oracle Solaris 11 Security* contains information about issues that programmers should pay attention to when coding network applications. The chapter also contains information on the interfaces provided by the Oracle Solaris operating system to help make your application more resilient and secure.

This manual also assumes basic competence in programming, a working familiarity with the C programming language, and familiarity with the UNIX operating system, particularly networking concepts. For more information on UNIX networking basics, see the following books:

- *Unix Network Programming, Volume 1: The Sockets Networking API (3rd Edition)* By W. Richard Stevens, Bill Fenner, & Andrew M. Rudoff.
- *UNIX Network Programming, Volume 2: Interprocess Communications (2nd Edition)* By W. Richard Stevens.

Organization of the Manual

The services and capabilities of the basic system interfaces and basic network interfaces of the Oracle Solaris OS platform are described in the following chapters.

[Chapter 1, “Memory and CPU Management,”](#) describes the interfaces that create and manage memory mappings, do high performance file I/O, and control other aspects of memory management.

[Chapter 2, “Session Description Protocol API,”](#) describes the API framework and library functions for the Solaris implementation of the Session Description Protocol (SDP).

[Chapter 3, “Process Scheduler,”](#) describes the operation of the SunOS process scheduler, modification of the scheduler's behavior, the scheduler's interactions with process management interfaces, and performance effects.

[Chapter 4, “Locality Group APIs,”](#) describes the interfaces used to control the behavior and structure of locality groups and resource priority for threads within those groups.

[Chapter 5, “Input/Output Interfaces,”](#) describes basic and old-style buffered file I/O and other elements of I/O.

[Chapter 6, “Interprocess Communication,”](#) describes older forms of non-networked interprocess communication.

[Chapter 7, “Socket Interfaces,”](#) describes the use of sockets, which are the basic mode of networked communication.

[Chapter 8, “Programming With XTI and TLI,”](#) describes the use of XTI and TLI to do transport-independent networked communication.

[Chapter 9, “Packet Filtering Hooks,”](#) describes interfaces for developing network solutions at the kernel level such as security (packet filtering and firewall) solutions and network address translation (NAT) solutions.

[Chapter 10, “Transport Selection and Name-to-Address Mapping,”](#) describes the network selection mechanisms used by applications to select a network transport and its configuration.

[Chapter 11, “Real-time Programming and Administration,”](#) describes real-time programming facilities in the SunOS environment and their use.

[Chapter 12, “The Oracle Solaris ABI and ABI Tools,”](#) describes the Oracle Solaris Application Binary Interface (ABI) and the tools used to verify an application's compliance with the Oracle Solaris ABI, `apptest` and `apptestrace`.

[Appendix A, “UNIX Domain Sockets,”](#) describes UNIX domain sockets.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Description	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-2 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$

TABLE P-2 Shell Prompts *(Continued)*

Shell	Prompt
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

Memory and CPU Management

This chapter describes an application developer's view of virtual memory and CPU management in the Oracle Solaris operating system.

- “[Memory Management Interfaces](#)” on page 15 describes interfaces and cache control.
- Library level dynamic memory allocation and debugging are described in “[Library-Level Dynamic Memory](#)” on page 17.
- “[Other Memory Control Interfaces](#)” on page 19 describes other memory control interfaces.
- “[CPU Performance Counters](#)” on page 20 describes the use of CPU Performance Counters (CPC).

Memory Management Interfaces

Applications use the virtual memory facilities through several sets of interfaces. This section summarizes these interfaces. This section also provides examples of the interfaces' use.

Creating and Using Mappings

`mmap(2)` establishes a mapping of a named file system object into a process address space. A named file system object can also be partially mapped into a process address space. This basic memory management interface is very simple. Use `open(2)` to open the file, then use `mmap(2)` to create the mapping with appropriate access and sharing options. Then, proceed with your application.

The mapping established by `mmap(2)` replaces any previous mappings for the specified address range.

The flags `MAP_SHARED` and `MAP_PRIVATE` specify the type of mapping. You must specify a mapping type. If the `MAP_SHARED` flag is set, write operations modify the mapped object. No further operations on the object are needed to make the change. If the `MAP_PRIVATE` flag is set,

the first write operation to the mapped area creates a copy of the page. All further write operations reference the copy. Only modified pages are copied.

A mapping type is retained across a `fork(2)`.

After you have established the mapping through `mmap(2)`, the file descriptor used in the call is no longer used. If you close the file, the mapping remains until `munmap(2)` undoes the mapping. Creating a new mapping replaces an existing mapping.

A mapped file can be shortened by a call to `truncate`. An attempt to access the area of the file that no longer exists causes a SIGBUS signal.

Mapping `/dev/zero` gives the calling program a block of zero-filled virtual memory. This can also be done by setting the `MAP_ANON` flag and the file descriptor variable, `filedes` to `-1`. The size of the block is specified in the call to `mmap(2)`.

Some devices or files are useful only when accessed by mapping. Frame buffer devices used to support bit-mapped displays are an example of this phenomenon. Display management algorithms are much simpler to implement when the algorithms operate directly on the addresses of the display.

Removing Mappings

`munmap(2)` removes all mappings of pages in the specified address range of the calling process. `munmap(2)` has no affect on the objects that were mapped.

Cache Control

The virtual memory system in SunOS is a cache system, in which processor memory buffers data from file system objects. Interfaces are provided to control or interrogate the status of the cache.

Using `mincore`

The `mincore(2)` interface determines the residency of the memory pages in the address space covered by mappings in the specified range. Because the status of a page can change after `mincore` checks the page but before `mincore` returns the data, returned information can be outdated. Only locked pages are guaranteed to remain in memory.

Using `mlock` and `munlock`

`mlock(3C)` causes the pages in the specified address range to be locked in physical memory. References to locked pages in this process or in other processes do not result in page faults that require an I/O operation. Because this I/O operation interferes with normal operation of virtual

memory, as well as slowing other processes, the use of `mlock` is limited to the superuser. The limit to the number of pages that can be locked in memory is dependent on system configuration. The call to `mlock` fails if this limit is exceeded.

`munlock` releases the locks on physical pages. If multiple `mlock` calls are made on an address range of a single mapping, a single `munlock` call releases the locks. However, if different mappings to the same pages are locked by `mlock`, the pages are not unlocked until the locks on all the mappings are released.

Removing a mapping also releases locks, either through being replaced with an `mmap(2)` operation or removed with `munmap(2)`.

The copy-on-write event that is associated with a `MAP_PRIVATE` mapping transfers a lock on the source page to the destination page. Thus locks on an address range that includes `MAP_PRIVATE` mappings are retained transparently along with the copy-on-write redirection. For a discussion of this redirection, see “[Creating and Using Mappings](#)” on page 15.

Using `mlockall` and `munlockall`

`mlockall(3C)` and `munlockall(3C)` are similar to `mlock` and `munlock`, but `mlockall` and `munlockall` operate on entire address spaces. `mlockall` sets locks on all pages in the address space and `munlockall` removes all locks on all pages in the address space, whether established by `mlock` or `mlockall`.

Using `msync`

`msync(3C)` causes all modified pages in the specified address range to be flushed to the objects mapped by those addresses. This command is similar to `fsync(3C)`, which operates on files.

Library-Level Dynamic Memory

Library-level dynamic memory allocation provides an easy-to-use interface to dynamic memory allocation.

Dynamic Memory Allocation

The most often used interfaces are:

- `malloc(3C)`
- `free(3C)`
- `calloc(3C)`
- `watchmalloc(3MALLOC)`

Other dynamic memory allocation interfaces are `memalign(3C)`, `valloc(3C)`, and `realloc(3C)`

- `malloc` returns a pointer to a block of memory at least as large as the amount of memory that is requested. The block is aligned to store any type of data.
- `free` returns the memory that is obtained from `malloc`, `calloc`, `realloc`, `memalign`, or `valloc` to system memory. Trying to free a block that was not reserved by a dynamic memory allocation interface is an error that can cause a process to crash.
- `calloc` returns a pointer to a block of memory that is initialized to zeros. Memory reserved by `calloc` can be returned to the system through either `watchmalloc` or `free`. The memory is allocated and aligned to contain an array of a specified number of elements of a specified size.
- `memalign` allocates a specified number of bytes on a specified alignment boundary. The alignment boundary must be a power of 2.
- `valloc` allocates a specified number of bytes that are aligned on a page boundary.
- `realloc` changes the size of the memory block allocated to a process. `realloc` can be used to increase or reduce the size of an allocated block of memory. `realloc` is the only way to shrink a memory allocation without causing a problem. The location in memory of the reallocated block might be changed, but the contents up to the point of the allocation size change remain the same.

Dynamic Memory Debugging

The Oracle Solaris Studio software included tools that are useful in finding and eliminating errors in dynamic memory use.

`dbx` is an interactive, source-level, command-line debugging tool. You can use it to run a program in a controlled manner and to inspect the state of a stopped program. `dbx` gives you complete control of the dynamic execution of a program, including collecting performance and memory usage data, monitoring memory access, and detecting memory leaks. `dbxtool` provides a graphical user interface for `dbx`. See [Oracle Solaris Studio 12.3: Debugging a Program With `dbx`](#) for detailed information.

The Run Time Checking (RTC) tool in the Oracle Solaris Studio software lets you automatically detect runtime errors, such as memory access errors and memory leak, in a native code application during the development phase. It also lets you monitor memory usage. You cannot use runtime checking on Java code. See [Chapter 9, “Using Runtime Checking,” in Oracle Solaris Studio 12.3: Debugging a Program With `dbx`](#) for details on using the RTC facility.

You can also use the advanced development tool *Memory Error Discovery Tool (Discover)* for detecting memory access error. See the [Oracle Solaris Studio 12.3 Discover and Uncover User's Guide](#) for detailed information.

Oracle Solaris Studio is available on as a package to download and install on the Oracle Solaris 11 OS. For more information, see the [Oracle Solaris Studio website](#).

Other Memory Control Interfaces

This section discusses additional memory control interfaces.

Using `sysconf`

`sysconf(3C)` returns system dependent sizes of memory pages and applications should use `getpagesizes(3C)` to find out which memory pages are available to a running process. For portability, applications should not embed any constants that specify the size of a page. Note that varying page sizes are not unusual, even among implementations of the same instruction set.

Using `mprotect`

`mprotect(2)` assigns the specified protection to all pages in the specified address range. The protection cannot exceed the permissions that are allowed on the underlying object.

Using `brk` and `sbrk`

A *break* is the greatest valid data address in the process image that is not in the stack. When a program starts executing, the break value is normally set by `execve(2)` to the greatest address defined by the program and its data storage.

Use `brk(2)` to set the break to a greater address. You can also use `sbrk(2)` to add an increment of storage to the data segment of a process. You can get the maximum possible size of the data segment by a call to `getrlimit(2)`.

```
caddr_t  
brk(caddr_t addr);
```

```
caddr_t  
sbrk(intptr_t incr);
```

`brk` identifies the lowest data segment location not used by the caller as *addr*. This location is rounded up to the next multiple of the system page size.

`sbrk`, the alternate interface, adds *incr* bytes to the caller data space and returns a pointer to the start of the new data area.

CPU Performance Counters

This section describes developer interfaces for use of CPU Performance counters (CPC). Oracle Solaris applications can use CPC independent of the underlying counter architecture.

API Additions to libcpc

This section covers recent additions to the [libcpc\(3LIB\)](#) library. Please see the `libcpc` man page for information on older interfaces.

Initialization Interfaces

An application preparing to use the CPC facility initializes the library with a call to the `cpc_open()` function. This function returns a `cpc_t *` parameter that is used by the other interfaces. The syntax for the `cpc_open()` function is as follows:

```
cpc_t*cpc_open(int ver);
```

The value of the `ver` parameter identifies the version of the interface that the application is using. The `cpc_open()` function fails if the underlying counters are inaccessible or unavailable.

Hardware Query Interfaces

```
uint_t cpc_npics(cpc_t *cpc);
uint_t cpc_caps(cpc_t *cpc);
void cpc_walk_events_all(cpc_t *cpc, void *arg,
    void (*action)(void *arg, const char *event));
void cpc_walk_events_pic(cpc_t *cpc, uint_t picno, void *arg,
    void (*action)(void *arg, uint_t picno, const char *event));
void cpc_walk_attrs(cpc_t *cpc, void *arg,
    void (*action)(void *arg, const char *attr));
```

The `cpc_npics()` function returns the number of physical counters on the underlying processor.

The `cpc_caps()` function returns a `uint_t` parameter whose value is the result of the bitwise inclusive-OR operation performed on the capabilities that the underlying processor supports. There are two capabilities. The `CPC_CAP_OVERFLOW_INTERRUPT` capability enables the processor to generate an interrupt when a counter overflows. The `CPC_CAP_OVERFLOW_PRECISE` capability enables the processor to determine which counter generates an overflow interrupt.

The kernel maintains a list of the events that the underlying processor supports. Different physical counters on a single chip do not have to use the same list of events. The `cpc_walk_events_all()` function calls the `action()` routine for each processor-supported event without regard to physical counter. The `cpc_walk_events_pic()` function calls the `action()` routine for each processor-supported event on a specific physical counter. Both of these functions pass the `arg` parameter uninterpreted from the caller to each invocation of the `action()` function.

The platform maintains a list of attributes that the underlying processor supports. These attributes enable access to advanced processor-specific features of the performance counters. The `cpc_walk_attrs()` function calls the action routine on each attribute name.

Configuration Interfaces

```
cpc_set_t *cpc_set_create(cpc_t *cpc);
int cpc_set_destroy(cpc_t *cpc, cpc_set_t *set);
int cpc_set_add_request(cpc_t *cpc, cpc_set_t *set, const char *event,
                      uint64_t preset, uint_t flags, uint_t nattrs,
                      const cpc_attr_t *attrs);
int cpc_set_request_preset(cpc_t *cpc, cpc_set_t *set, int index,
                          uint64_t preset);
```

The opaque data type `cpc_set_t` represents collections of requests. The collections are called sets. The `cpc_set_create()` function creates an empty set. The `cpc_set_destroy()` function destroys a set and frees all the memory used by the set. Destroying a set releases the hardware resources the set uses.

The `cpc_set_add_request()` function adds requests to a set. The following list describes the parameters of a request.

<code>event</code>	A string that specifies the name of the event to count.
<code>preset</code>	A 64-bit unsigned integer that is used as the initial value of the counter.
<code>flags</code>	The results of the logical OR operation applied to a group of request flags.
<code>nattrs</code>	The number of attributes in the array that <code>attrs</code> points to.
<code>attrs</code>	A pointer to an array of <code>cpc_attr_t</code> structures.

The following list describes the valid request flags.

<code>CPC_COUNT_USER</code>	This flag enables counting of events that occur while the CPU is executing in user mode.
<code>CPC_COUNT_SYSTEM</code>	This flag enables counting of events that occur while the CPU is executing in privileged mode.
<code>CPC_OVF_NOTIFY_EMT</code>	This flag requests notification of hardware counter overflow.

The CPC interfaces pass attributes as an array of `cpc_attr_t` structures.

When the `cpc_set_add_request()` function returns successfully, it returns an index. The index references the data generated by the request added by the call to the `cpc_set_add_request()` function.

The `cpc_set_request_preset()` function changes the preset value of a request. This enables the re-binding of an overflowed set with new presets.

The `cpc_walk_requests()` function calls a user-provided `action()` routine on each request in `cpc_set_t`. The value of the `arg` parameter is passed to the user routine without interpretation. The `cpc_walk_requests()` function allows applications to print the configuration of each request in a set. The syntax for the `cpc_walk_requests()` function is as follows:

```
void cpc_walk_requests(cpc_t *cpc, cpc_set_t *set, void *arg,
void (*action)(void *arg, int index, const char *event,
uint64_t preset, uint_t flags, int nattrs,
const cpc_attr_t *attrs));
```

Binding

The interfaces in this section bind the requests in a set to the physical hardware and set the counters to a starting position.

```
int cpc_bind_curlwp(cpc_t *cpc, cpc_set_t *set, uint_t flags);
int cpc_bind_pctx(cpc_t *cpc, pctx_t *pctx, id_t id, cpc_set_t *set,
uint_t flags);
int cpc_bind_cpu(cpc_t *cpc, processorid_t id, cpc_set_t *set,
uint_t flags);
int cpc_unbind(cpc_t *cpc, cpc_set_t *set);
```

The `cpc_bind_curlwp()` function binds the set to the calling LWP. The set's counters are virtualized to this LWP and count the events that occur on the CPU while the calling LWP runs. The only flag that is valid for the `cpc_bind_curlwp()` routine is `CPC_BIND_LWP_INHERIT`.

The `cpc_bind_pctx()` function binds the set to a LWP in a process that is captured with [libpctx\(3LIB\)](#). This function has no valid flags.

The `cpc_bind_cpu()` function binds the set to the processor specified in the `id` parameter. Binding a set to a CPU invalidates existing performance counter contexts on the system. This function has no valid flags.

The `cpc_unbind()` function stops the performance counters and releases the hardware that is associated with the bound set. If a set is bound to a CPU, the `cpc_unbind()` function unbinds the LWP from the CPU and releases the CPC pseudo-device.

Sampling

The interfaces described in this section enable the return of data from the counters to the application. Counter data resides in an opaque data structure named `cpc_buf_t`. This data structure takes a snapshot of the state of counters in use by a bound set and includes the following information:

- The 64-bit values of each counter
- The timestamp of the most recent hardware snapshot

- A cumulative CPU cycle counter that counts the number of CPU cycles the processor has used on the bound set

```
cpc_buf_t *cpc_buf_create(cpc_t *cpc, cpc_set_t *set);
int cpc_buf_destroy(cpc_t *cpc, cpc_buf_t *buf);
int cpc_set_sample(cpc_t *cpc, cpc_set_t *set, cpc_buf_t *buf);
```

The `cpc_buf_create()` function creates a buffer that stores data from the set specified in `cpc_set_t`. The `cpc_buf_destroy()` function frees the memory that is associated with the given `cpc_buf_t`. The `cpc_buf_sample()` function takes a snapshot of the counters that are counting on behalf of the specified set. The specified set must already be bound and have a buffer created before calling the `cpc_buf_sample()` function.

Sampling into a buffer does not update the preset of the requests that are associated with that set. When a buffer is sampled with the `cpc_buf_sample()` function, then unbound and bound again, counts start from the request's preset as in the original call to the `cpc_set_add_request()` function.

Buffer Operations

The following routines provide access to the data in a `cpc_buf_t` structure.

```
int cpc_buf_get(cpc_t *cpc, cpc_buf_t *buf, int index, uint64_t *val);
int cpc_buf_set(cpc_t *cpc, cpc_buf_t *buf, int index, uint64_t *val);
hrtime_t cpc_buf_hrtime(cpc_t *cpc, cpc_buf_t *buf);
uint64_t cpc_buf_tick(cpc_t *cpc, cpc_buf_t *buf);
int cpc_buf_sub(cpc_t *cpc, cpc_buf_t *result, cpc_buf_t *left,
               cpc_buf_t *right);
int cpc_buf_add(cpc_t *cpc, cpc_buf_t *result, cpc_buf_t *left,
               cpc_buf_t *right);
int cpc_buf_copy(cpc_t *cpc, cpc_buf_t *dest, cpc_buf_t *src);
void cpc_buf_zero(cpc_t *cpc, cpc_buf_t *buf);
```

The `cpc_buf_get()` function retrieves the value of the counter that is identified by the *index* parameter. The *index* parameter is a value that is returned by the `cpc_set_add_request()` function before the set is bound. The `cpc_buf_get()` function stores the value of the counter at the location indicated by the *val* parameter.

The `cpc_buf_set()` function sets the value of the counter that is identified by the *index* parameter. The *index* parameter is a value that is returned by the `cpc_set_add_request()` function before the set is bound. The `cpc_buf_set()` function sets the counter's value to the value at the location indicated by the *val* parameter. Neither the `cpc_buf_get()` function nor the `cpc_buf_set()` function change the preset of the corresponding CPC request.

The `cpc_buf_hrtime()` function returns the high resolution timestamp that indicates when the hardware was sampled. The `cpc_buf_tick()` function returns the number of CPU clock cycles that have elapsed while the LWP is running.

The `cpc_buf_sub()` function computes the difference between the counters and tick values that are specified in the *left* and *right* parameters. The `cpc_buf_sub()` function stores the results in *result*. A given invocation of the `cpc_buf_sub()` function must have all `cpc_buf_t` values originate from the same `cpc_set_t` structure. The *result* index contains the result of the *left* - *right* computation for each request index in the buffers. The result index also contains the *tick* difference. The `cpc_buf_sub()` function sets the high-resolution timestamp of the destination buffer to the most recent time of the *left* or *right* buffers.

The `cpc_buf_add()` function computes the total of the counters and tick values that are specified in the *left* and *right* parameters. The `cpc_buf_add()` function stores the results in *result*. A given invocation of the `cpc_buf_add()` function must have all `cpc_buf_t` values originate from the same `cpc_set_t` structure. The *result* index contains the result of the *left* + *right* computation for each request index in the buffers. The result index also contains the *tick* total. The `cpc_buf_add()` function sets the high-resolution timestamp of the destination buffer to the most recent time of the *left* or *right* buffers.

The `cpc_buf_copy()` function makes *dest* identical to *src*.

The `cpc_buf_zero()` function sets everything in *buf* to zero.

Activation Interfaces

This section describes activation interfaces for CPC.

```
int cpc_enable(cpc_t *cpc);
int cpc_disable(cpc_t *cpc);
```

These two interfaces respectively enable and disable counters of any set that is bound to the executing LWP. Use of these interfaces enables an application to designate code of interest while deferring the counter configuration to a controlling process by using `libpctx`.

Error Handling Interfaces

This section describes CPC's error handling interfaces.

```
typedef void (cpc_errhdlr_t)(const char *fn, int subcode, const char *fmt,
                             va_list ap);
void cpc_seterrhdlr(cpc_t *cpc, cpc_errhdlr_t *errhdlr);
```

These two interfaces allow the passage of a `cpc_t` handle. The `cpc_errhdlr_t` handle takes an integer subcode in addition to a string. The integer *subcode* describes the specific error that was encountered by the function that the *fn* argument refers to. The integer *subcode* simplifies an application's recognition of error conditions. The string value of the *fmt* argument contains an internationalized description of the error subcode and is suitable for printing.

Session Description Protocol API

The Session Description Protocol (SDP) describes multimedia sessions. The SDP API discussed in this chapter contains function calls you can use to add SDP functionality to your applications.

Session Description API Overview

The function calls that make up the SDP API are provided by the shared object `libcommputil.so.1`. The functions in this shared object parse the SDP description and check the description's syntax.

The `sdp.h` header file defines the `sdp_session_t` structure, which contains the following members:

```
typedef struct sdp_session {
    int          sdp_session_version; /* SDP session version */
    int          s_version;           /* SDP version field */
    sdp_origin_t *s_origin;           /* SDP origin field */
    char         *s_name;              /* SDP name field */
    char         *s_info;              /* SDP info field */
    char         *s_uri;               /* SDP uri field */
    sdp_list_t   *s_email;             /* SDP email field */
    sdp_list_t   *s_phone;             /* SDP phone field */
    sdp_conn_t   *s_conn;              /* SDP connection field */
    sdp_bandwidth_t *s_bw;             /* SDP bandwidth field */
    sdp_time_t   *s_time;              /* SDP time field */
    sdp_zone_t   *s_zone;              /* SDP zone field */
    sdp_key_t    *s_key;               /* SDP key field */
    sdp_attr_t   *s_attr;              /* SDP attribute field */
    sdp_media_t  *s_media;             /* SDP media field */
} sdp_session_t;
```

The `sdp_session_version` member tracks the version of the structure. The initial value of the `sdp_session_version` member is `SDP_SESSION_VERSION_1`.

The `sdp_origin_t` structure contains the following members:

```

typedef struct sdp_origin {
    char      *o_username; /* username of the originating host */
    uint64_t  o_id;        /* session id */
    uint64_t  o_version;   /* version number of this session */
                                /* description */
    char      *o_nettype;  /* type of network */
    char      *o_addrtype; /* type of the address */
    char      *o_address;  /* address of the machine from which */
                                /* session was created */
} sdp_origin_t;

```

The `sdp_conn_t` structure contains the following members:

```

typedef struct sdp_conn {
    char      *c_nettype;   /* type of network */
    char      *c_addrtype; /* type of the address */
    char      *c_address;  /* unicast-address or multicast */
                                /* address */
    int       c_addrcount; /* number of addresses (case of */
                                /* multicast address with layered */
                                /* encodings */
    struct    sdp_conn *c_next; /* pointer to next connection */
                                /* structure; there could be several */
                                /* connection fields in SDP description */
    uint8_t  c_ttl;        /* TTL value for IPV4 multicast address */
} sdp_conn_t;

```

The `sdp_bandwidth_t` structure contains the following members:

```

typedef struct sdp_bandwidth {
    char      *b_type; /* info needed to interpret b_value */
    uint64_t  b_value; /* bandwidth value */
    struct    sdp_bandwidth *b_next; /* pointer to next bandwidth structure*/
                                /* (there could be several bandwidth */
                                /* fields in SDP description */
} sdp_bandwidth_t;

```

The `sdp_list_t` structure is a linked list of void pointers. This structure holds SDP fields. In the case of SDP structure fields such as email and phone, the void pointers point to character buffers. Use this structure to hold information in cases where the number of elements is not predefined, as in the case of repeated offset fields, where the void pointer holds integer values.

The `sdp_list_t` structure contains the following members:

```

typedef struct sdp_list {
    void      *value; /* string values in case of email, phone and */
                                /* format (in media field) or integer values */
                                /* in case of offset (in repeat field) */
    struct    sdp_list *next; /* pointer to the next node in the list */
} sdp_list_t;

```

The `sdp_repeat_t` structure will always be part of the time structure `sdp_time_t`. The repeat field does not appear alone in SDP descriptions and is always associated with the time field.

The `sdp_repeat_t` structure contains the following members:

```

typedef struct sdp_repeat {
    uint64_t      r_interval; /* repeat interval, e.g. 86400 seconds */
                        /* (1 day) */
    uint64_t      r_duration; /* duration of session, e.g. 3600 */
                        /* seconds (1 hour) */
    sdp_list_t    *r_offset; /* linked list of offset values; each */
                        /* represents offset from start-time */
                        /* in the SDP time field */
    struct sdp_repeat *r_next; /* pointer to next repeat structure; */
                        /* there could be several repeat */
                        /* fields in the SDP description */
}

```

The `sdp_time_t` structure contains the following members:

```

typedef struct sdp_time {
    uint64_t      t_start; /* start-time for a session */
    uint64_t      t_stop; /* end-time for a session */
    sdp_repeat_t  *t_repeat; /* points to the SDP repeat field */
    struct sdp_time *t_next; /* pointer to next time field; there */
                        /* could there could be several time */
                        /* fields in SDP description */
} sdp_time_t;

```

The `sdp_zone_t` structure contains the following members:

```

typedef struct sdp_zone {
    uint64_t      z_time; /* base time */
    char          *z_offset; /* offset added to z_time to determine */
                        /* session time; mainly used for daylight */
                        /* saving time conversions */
    struct sdp_zone *z_next; /* pointer to next zone field; there */
                        /* could be several <adjustment-time> */
                        /* <offset> pairs within a zone field */
} sdp_zone_t;

```

The `sdp_key_t` structure contains the following members:

```

typedef struct sdp_key {
    char *k_method; /* key type */
    char *k_enckey; /* encryption key */
} sdp_key_t;

```

The `sdp_attr_t` structure contains the following members:

```

typedef struct sdp_attr {
    char *a_name; /* name of the attribute */
    char *a_value; /* value of the attribute */
    struct sdp_attr *a_next; /* pointer to the next attribute */
                        /* structure; there could be several */
                        /* attribute fields within SDP description */
} sdp_attr_t;

```

The `sdp_media_t` structure contains the following members:

```

typedef struct sdp_media {
    char *m_name; /* name of the media such as "audio", */
                        /* "video", "message" */
}

```

```
uint_t      m_port;      /* transport layer port information */
int         m_portcount; /* number of ports in case of */
            /* hierarchically encoded streams */
char        *m_proto;    /* transport protocol */
sdp_list_t  *m_format;   /* media format description */
char        *m_info;     /* media info field */
sdp_conn_t  *m_conn;     /* media connection field */
sdp_bandwidth_t *m_bw;   /* media bandwidth field */
sdp_key_t   *m_key;      /* media key field */
sdp_attr_t  *m_attr;     /* media attribute field */
struct sdp_media *m_next; /* pointer to next media structure; */
            /* there could be several media */
            /* sections in SDP description */
sdp_session_t *m_session; /* pointer to the session structure */
} sdp_media_t;
```

SDP Library Functions

The API library functions support the following operations:

- Creating the SDP session structure
- Searching within the SDP session structure
- Shutting down an SDP session structure
- Utility functions

Creating the SDP Session Structure

The first step in creating a new SDP session structure is allocating memory for the new structure by calling the `sdp_new_session()` function. This function returns a pointer to the new session structure. The other functions in this section use that pointer to construct the new session structure. Once you are done constructing the new session structure, convert it to a string representation with the `sdp_session_to_str()` function.

Creating a New SDP Session Structure

```
sdp_session_t *sdp_new_session();
```

The `sdp_new_session()` function allocates memory for a new SDP session structure that is specified by the `session` parameter and assigns a version number to that new structure. You can free the memory that is allocated to the session structure by calling the `sdp_free_session()` function.

Return Values: The `sdp_new_session()` function returns the newly allocated SDP session structure when the function completes successfully. The function returns `NULL` in the case of failure.

Adding an Origin Field to the SDP Session Structure

```
int sdp_add_origin(sdp_session_t *session, const char *name, uint64_t id,
uint64_t ver, const char *nettype, const char *addrtype, const char *address);
```

The `sdp_add_origin()` function adds the ORIGIN (o=) SDP field to the session structure that is specified by the value of the session parameter (`sdp_session_t`) using the `name`, `id`, `ver`, `nettype`, `addrtype`, and `address` parameters.

Return Values: The `sdp_add_origin()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Name Field to the SDP Session Structure

```
int sdp_add_name(sdp_session_t *session, const char *name);
```

The `sdp_add_name()` function adds the NAME (s=) SDP field to the session structure that is specified by the value of the session parameter (`sdp_session_t`) using the `name` parameter.

Return Values: The `sdp_add_name()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding an Information Field to the SDP Session Structure

```
int sdp_add_information(char **information, const char *value);
```

The `sdp_add_information()` function adds the INFO (i=) SDP field to the session structure (`sdp_session_t`) or the media structure (`sdp_media_t`) using the `value` parameter. This field can go into the media or the session section of an SDP description. You must pass either `&session->s_info` or `&media->m_info` as the first argument to specify the section.

Return Values: The `sdp_add_information()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a URI Field to the SDP Session Structure

```
int sdp_add_uri(sdp_session_t *session, const char *uri);
```

The `sdp_add_uri()` function adds the URI (u=) SDP field to the session structure that is specified by the value of the session parameter (`sdp_session_t`) using the `uri` parameter.

Return Values: The `sdp_add_uri()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding an Email Field to the SDP Session Structure

```
int sdp_add_email(sdp_session_t *session, const char *email);
```

The `sdp_add_email()` function adds the `EMAIL` (e=) SDP field to the session structure that is specified by the value of the `session` parameter (`sdp_session_t`) using the `email` parameter.

Return Values: The `sdp_add_email()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Telephone Field to the SDP Session Structure

```
int sdp_add_phone(sdp_session_t *session, const char *email);
```

The `sdp_add_phone()` function adds the `PHONE` (p=) SDP field to the session structure that is specified by the value of the `session` parameter (`sdp_session_t`) using the `phone` parameter.

Return Values: The `sdp_add_phone()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Connection Field to the SDP Session Structure

```
int sdp_add_connection(sdp_conn_t **conn, const char *nettype, const char *addrtype, const char *address, uint8_t ttl, int addrcount);
```

The `sdp_add_connection()` function adds the `CONNECTION` (c=) SDP field to either the session structure (`sdp_session_t`) or the media structure (`sdp_media_t`) using the `nettype`, `addrtype`, `address`, `ttl`, and `addrcount` parameters. For IPv4 or IPv6 unicast addresses, set the values of the `ttl` and `addrcount` parameters to zero. For multicast addresses, set the value of the `ttl` parameter between zero and 255. A multicast address cannot have an `addrcount` parameter with a value of zero.

This field can go into the media or the session section of an SDP description. You must pass either `&session->s_info` or `&media->m_info` as the first argument to specify the section.

Return Values: The `sdp_add_connection()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Bandwidth Field to the SDP Session Structure

```
int sdp_add_bandwidth(sdp_bandwidth_t **bw, const char *type, uint64_t value);
```

The `sdp_add_bandwidth()` function adds the BANDWIDTH (b=) SDP field to either the session structure (`sdp_session_t`) or the media structure (`sdp_media_t`) using the `type` and `value` parameters.

This field can go into the media or the session section of an SDP description. You must pass either `&session->s_info` or `&media->m_info` as the first argument to specify the section.

Return Values: The `sdp_add_bandwidth()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Time Field to the SDP Session Structure

```
int sdp_add_time(sdp_session_t *session, uint64_t starttime, uint64_t stoptime,
sdp_time_t **time);
```

The `sdp_add_time()` function adds the TIME (t=) SDP field to the session structure using the values of the `starttime` and `stoptime` parameters. This function creates a new time structure and returns the pointer to that structure in the `time` parameter.

Return Values: The `sdp_add_time()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Repeat Field to the SDP Session Structure

```
int sdp_add_repeat(sdp_time_t *time, uint64_t interval, uint64_t duration, const
char *offset);
```

The `sdp_add_repeat()` function adds the REPEAT (r=) SDP field to the session structure using the values of the `interval`, `duration`, and `offset` parameters. The value of the `offset` parameter is a string that holds one or more offset values, such as `60` or `60 1d 3h`. The value of the `time` parameter is the pointer to the time structure that the `sdp_add_time()` function creates.

Return Values: The `sdp_add_repeat()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Zone Field to the SDP Session Structure

```
int sdp_add_zone(sdp_session_t *session, uint64_t time, const char *offset);
```

The `sdp_add_zone()` function adds the ZONE (z=) SDP field to the session structure that is specified by the value of the session parameter (`sdp_session_t`) using the *time* and *offset* parameters. You can add multiple time and offset values for a single zone field by calling this function for each time/offset value pair.

Return Values: The `sdp_add_zone()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Key Field to the SDP Session Structure

```
int sdp_add_key(sdp_key_t **key, const char *method, const char *enckey);
```

The `sdp_add_key()` function adds the KEY (k=) SDP field to the session structure (`sdp_session_t`) or the media structure (`sdp_media_t`) using the *method* and *enckey* parameters. This field can go into the media or the session section of an SDP description. You must pass either `&session->s_info` or `&media->m_info` as the first argument to specify the section.

Return Values: The `sdp_add_key()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding an Attribute Field to the SDP Session Structure

```
int sdp_add_attribute(sdp_attr_t **attr, const char *name, const char *value);
```

The `sdp_add_attribute()` function adds the ATTRIBUTE (a=) SDP field to the session structure (`sdp_session_t`) or the media structure (`sdp_media_t`) using the *name* and *value* parameters. This field can go into the media or the session section of an SDP description. You must pass either `&session->s_info` or `&media->m_info` as the first argument to specify the section.

Return Values: The `sdp_add_attribute()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Adding a Media Field to the SDP Session Structure

```
int sdp_add_media(sdp_session_t *session, const char *name, uint_t port, int portcount, const char *protocol, const char *format, sdp_media_t **media);
```

The `sdp_add_media()` function adds the MEDIA (m=) SDP field to the session structure that is specified by the value of the session parameter (`sdp_session_t`) using the values of the *name*, *port*, *portcount*, *protocol*, and *format* parameters. The *format* parameter is a string that holds one or more values, such as the string `0 32 97`.

This function creates a new media structure and returns a pointer to that structure in the *media* parameter. Functions that add SDP fields to the media structure use this pointer.

Return Values: The `sdp_add_media()` function returns 0 when the function completes successfully. When mandatory parameters are not present, the function returns `EINVAL`. When memory allocation fails, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

Code Sample: Building an SDP Session Structure

This example uses the functions in this section to create a new SDP session structure, add fields to the structure, and convert a finished structure to its string representation. At the end of the example, the program calls the `sdp_free_session()` function to free the session.

EXAMPLE 2-1 Building an SDP Session Structure

```

/* SDP Message we will be building
"v=0\r\n\
o=Alice 2890844526 2890842807 IN IP4 10.47.16.5\r\n\
s=-\r\n\
i=A Seminar on the session description protocol\r\n\
u=http://www.example.com/seminars/sdp.pdf\r\n\
e=alice@example.com (Alice Smith)\r\n\
p=+1 911-345-1160\r\n\
c=IN IP4 10.47.16.5\r\n\
b=CT:1024\r\n\
t=2854678930 2854679000\r\n\
r=604800 3600 0 90000\r\n\
z=2882844526 -1h 2898848070 0h\r\n\
a=recvonly\r\n\
m=audio 49170 RTP/AVP 0\r\n\
i=audio media\r\n\
b=CT:1000\r\n\
k=prompt\r\n\
m=video 51372 RTP/AVP 99 90\r\n\
i=video media\r\n\
a=rtpmap:99 h232-199/90000\r\n\
a=rtpmap:90 h263-1998/90000\r\n\
*/

#include <stdio.h>
#include <string.h>
#include <errno.h>
#include <sdp.h>

int main ()
{
    sdp_session_t *my_sess;
    sdp_media_t *my_media;
    sdp_time_t *my_time;
    char *b_sdp;

    my_sess = sdp_new_session();
    if (my_sess == NULL) {
return (ENOMEM);

```

EXAMPLE 2-1 Building an SDP Session Structure (Continued)

```

    }
    my_sess->version = 0;
    if (sdp_add_name(my_sess, "-") != 0)
goto err_ret;
    if (sdp_add_origin(my_sess, "Alice", 2890844526ULL, 2890842807ULL,
    "IN", "IP4", "10.47.16.5") != 0)
goto err_ret;
    if (sdp_add_information(&my_sess->s_info, "A Seminar on the session"
    "description protocol") != 0)
goto err_ret;
    if (sdp_add_uri(my_sess, "http://www.example.com/seminars/sdp.pdf")
    != 0)
goto err_ret;
    if (sdp_add_email(my_sess, "alice@example.com (Alice smith)") != 0)
goto err_ret;
    if (sdp_add_phone(my_sess, "+1 911-345-1160") != 0)
goto err_ret;
    if (sdp_add_connection(&my_sess->s_conn, "IN", "IP4", "10.47.16.5",
    0, 0) != 0)
goto err_ret;
    if (sdp_add_bandwidth(&my_sess->s_bw, "CT", 1024) != 0)
goto err_ret;
    if (sdp_add_time(my_sess, 2854678930ULL, 2854679000ULL, &my_time)
    != 0)
goto err_ret;
    if (sdp_add_repeat(my_time, 604800ULL, 3600ULL, "0 90000") != 0)
goto err_ret;
    if (sdp_add_zone(my_sess, 2882844526ULL, "-1h") != 0)
goto err_ret;
    if (sdp_add_zone(my_sess, 2898848070ULL, "0h") != 0)
goto err_ret;
    if (sdp_add_attribute(&my_sess->s_attr, "sendrecv", NULL) != 0)
goto err_ret;
    if (sdp_add_media(my_sess, "audio", 49170, 1, "RTP/AVP",
    "0", &my_media) != 0)
goto err_ret;
    if (sdp_add_information(&my_media->m_info, "audio media") != 0)
goto err_ret;
    if (sdp_add_bandwidth(&my_media->m_bw, "CT", 1000) != 0)
goto err_ret;
    if (sdp_add_key(&my_media->m_key, "prompt", NULL) != 0)
goto err_ret;
    if (sdp_add_media(my_sess, "video", 51732, 1, "RTP/AVP",
    "99 90", &my_media) != 0)
goto err_ret;
    if (sdp_add_information(&my_media->m_info, "video media") != 0)
goto err_ret;
    if (sdp_add_attribute(&my_media->m_attr, "rtpmap",
    "99 h232-199/90000") != 0)
goto err_ret;
    if (sdp_add_attribute(&my_media->m_attr, "rtpmap",
    "90 h263-1998/90000") != 0)
goto err_ret;
    b_sdp = sdp_session_to_str(my_sess, &error);
    /*

```

EXAMPLE 2-1 Building an SDP Session Structure (Continued)

```

    * b_sdp is the string representation of my_sess structure
    */

    free(b_sdp);
    sdp_free_session(my_sess);
    return (0);
err_ret:
    free(b_sdp);
    sdp_free_session(my_sess);
    return (1);
}

```

Searching the SDP Session Structure

The functions in this section search the SDP session structure for specific values and return pointers to those values.

Finding an Attribute in an SDP Session Structure

```
sdp_attr_t *sdp_find_attribute(sdp_attr_t *attr, const char *name);
```

The `sdp_find_attribute()` function searches the attribute list that is specified by the `attr` parameter for the attribute name that is specified by the `name` parameter.

Return Values: The `sdp_find_attribute()` function returns a pointer to the attribute (`sdp_attr_t *`) that is specified by the `name` parameter when the function completes successfully. In all other cases, the `sdp_find_attribute()` function returns a value of `NULL`.

EXAMPLE 2-2 Using the `sdp_find_attribute()` Function

The incomplete SDP description in this example has an audio section.

```

m=audio 49170 RTP/AVP 0 8
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=sendonly
a=ptime:10000
a=maxptime:20000

/*
 * Assuming that above description is parsed using sdp_parse and that
 * the parsed structure is in "session" sdp_session_t structure.
 */

sdp_attr_t *ptime;
sdp_attr_t *max_ptime;
sdp_media_t *media = session->s_media;

if ((ptime = sdp_find_attribute(media->m_attr, "ptime")) == NULL)
/* ptime attribute not present */

```

EXAMPLE 2-2 Using the `sdp_find_attribute()` Function (Continued)

```
else if((max_ptime = sdp_find_attribute(media->m_attr,
    "maxptime")) == NULL)
/* max_ptime attribute not present */
```

Finding Media in an SDP Session Structure

```
sdp_media_t *sdp_find_media(sdp_media_t *media, const char *name);
```

The `sdp_find_media()` function searches the media list that is specified by the *media* parameter for the media entry that is specified by the *name* parameter.

Return Values: The `sdp_find_media()` function returns a pointer to the media list entry (`sdp_media_t *`) that is specified by the *name* parameter when the function completes successfully. In all other cases, the `sdp_find_media()` function returns a value of `NULL`.

EXAMPLE 2-3 Using the `sdp_find_media()` Function

The incomplete SDP description in this example has two sections, an audio section and a video section.

```
m=audio 49170 RTP/AVP 0 8
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
m=video 51372 RTP/AVP 31 32
a=rtpmap:31 H261/90000
a=rtpmap:32 MPV/90000

/*
 * Assuming that above description is parsed using sdp_parse() and that
 * the parsed structure is in "session" sdp_session_t structure.
 */

sdp_media_t      *my_media;
my_media = sdp_find_media(session->s_media, "video");

/*
 * my_media now points to the structure containg video media section
 * information
 */
```

Finding a Media Format in an SDP Session Structure

```
sdp_attr_t *sdp_find_media_rtpmap(sdp_media_t *media, const char *format);
```

The `sdp_find_media_rtpmap()` function searches the attribute list of the media structure that is specified by the *media* parameter for the format entry that is specified by the *format* parameter.

Return Values: The `sdp_find_media_rtpmap()` function returns a pointer to the format entry (`sdp_attr_t *`) that is specified by the *name* parameter when the function completes successfully. In all other cases, the `sdp_find_media()` function returns a value of `NULL`.

EXAMPLE 2-4 Using the `sdp_find_media_rtpmap()` Function

The incomplete SDP description in this example has two sections, an audio section and a video section.

```

m=audio 49170 RTP/AVP 0 8
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
m=video 51372 RTP/AVP 31 32
a=rtpmap:31 H261/90000
a=rtpmap:32 MPV/90000

/*
 * Assuming that above description is parsed using sdp_parse() and that
 * the parsed structure is in "session" sdp_session_t structure.
 */

sdp_media_t      *video;
sdp_attr_t       *mpv;

video = sdp_find_media(session->s_media, "video");
mpv = sdp_find_media_rtpmap(video, "32");

/*
 * Now the attribute structure sdp_attr_t, mpv will be having
 * values from the attribute field "a=rtpmap:32 MPV/90000"
 */

```

Shutting Down the SDP Session Structure

The functions in this section perform the following functions:

- Removing fields from an SDP session structure
- Freeing an SDP session structure

Deleting Fields From the SDP Session Structure

```
int sdp_delete_all_field(sdp_session_t *session, const char field);
```

The `sdp_delete_all_field()` function deletes all occurrences of the SDP field that is specified by the *field* parameter from the SDP structure. For example, if an SDP structure has three BANDWIDTH (b=) fields, calling this function with a value of `SDP_BANDWIDTH_FIELD` in the *field* parameter deletes all three BANDWIDTH fields from the session structure.

Return Values: The `sdp_delete_all_field()` function returns 0 when the function completes successfully. When the session argument is NULL or the field type is unknown, the function returns EINVAL. The value of `errno` does not change in the event of an error.

Deleting Fields From the SDP Media Structure

```
int sdp_delete_all_media_field(sdp_media_t *media, const char field);
```

The `sdp_delete_all_media_field()` function deletes all occurrences of the SDP field that is specified by the *field* parameter from the SDP media structure.

Return Values: The `sdp_delete_all_media_field()` function returns 0 when the function completes successfully. When the session argument is NULL or the field type is unknown, the function returns EINVAL. The value of `errno` does not change in the event of an error.

Deleting Media From the SDP Media Structure

```
int sdp_delete_media(sdp_media_t **l_media, sdp_media_t *media);
```

The `sdp_delete_media()` function deletes the media entry specified by the *media* parameter from the media list. This function finds the specified media entry by calling the `sdp_find_media()` function. This function frees the memory that is allocated to the media structure after deleting the media entry.

Return Values: The `sdp_delete_media()` function returns 0 when the function completes successfully. When the session argument is NULL or mandatory arguments do not exist, the function returns EINVAL. The value of `errno` does not change in the event of an error.

Deleting an Attribute From the SDP Media Structure

```
int sdp_delete_attribute(sdp_attr_t **l_attr, sdp_attr_t *attr);
```

The `sdp_delete_attribute()` function deletes the attribute specified by the *attr* parameter from the media list. This function finds the specified attribute by calling either the `sdp_find_media_rtpmap()` function or the `sdp_find_attribute()` function. This function frees the memory that is allocated to the attribute structure after deleting the attribute.

Return Values: The `sdp_delete_attribute()` function returns 0 when the function completes successfully. When the session argument is NULL or mandatory arguments do not exist, the function returns EINVAL. The value of `errno` does not change in the event of an error.

Deleting an Attribute From the SDP Media Structure

```
void sdp_free_session(sdp_session_t *session);
```

The `sdp_free_session()` function destroys the session specified by the *session* parameter and frees the resources that are associated with that structure.

SDP API Utility Functions

The functions in this section parse and populate the SDP session structure, clone an existing session, and convert an existing session to a string representation.

Parsing the SDP Session Structure

```
int sdp_parse(const char *sdp_info, int len, int flags, sdp_session_t **session,
uint_t *p_error);
```

The `sdp_parse()` function parses the SDP description in the `sdp_info` parameter and populates the `sdp_session_t` structure. The `len` parameter specifies the length of the character buffer `sdp_info`. The function allocates the memory required for the `sdp_session_t` structure. To free that memory, call the `sdp_free_session()` function.

The value of the `flags` parameter must be set to zero. When the `flags` parameter has a nonzero value, the `sdp_parse()` function fails with a return value of `EINVAL` and sets the value of `*session` to `NULL`.

The `p_error` parameter takes on the values of any fields that have parsing errors. This parameter cannot have a value of `NULL`. Possible values for the `p_error` parameter are in the following list:

<code>SDP_VERSION_ERROR</code>	<code>0x00000001</code>
<code>SDP_ORIGIN_ERROR</code>	<code>0x00000002</code>
<code>SDP_NAME_ERROR</code>	<code>0x00000004</code>
<code>SDP_INFO_ERROR</code>	<code>0x00000008</code>
<code>SDP_URI_ERROR</code>	<code>0x00000010</code>
<code>SDP_EMAIL_ERROR</code>	<code>0x00000020</code>
<code>SDP_PHONE_ERROR</code>	<code>0x00000040</code>
<code>SDP_CONNECTION_ERROR</code>	<code>0x00000080</code>
<code>SDP_BANDWIDTH_ERROR</code>	<code>0x00000100</code>
<code>SDP_TIME_ERROR</code>	<code>0x00000200</code>
<code>SDP_REPEAT_TIME_ERROR</code>	<code>0x00000400</code>
<code>SDP_ZONE_ERROR</code>	<code>0x00000800</code>
<code>SDP_KEY_ERROR</code>	<code>0x00001000</code>
<code>SDP_ATTRIBUTE_ERROR</code>	<code>0x00002000</code>
<code>SDP_MEDIA_ERROR</code>	<code>0x00004000</code>
<code>SDP_FIELDS_ORDER_ERROR</code>	<code>0x00008000</code>
<code>SDP_MISSING_FIELDS</code>	<code>0x00010000</code>

When the SDP structure violates RFC 4566 by having fields out of order, the `sdp_parse()` function sets the value of the `p_error` parameter to `SDP_FIELDS_ORDER_ERROR`. When the SDP structure violates RFC 4566 by lacking mandatory fields, the `sdp_parse()` function sets the value of the `p_error` parameter to `SDP_MISSING_FIELDS`.

The `sdp_parse()` function continues to parse after processing a field with a parsing error, but the field with the parsing error will not be present in the resulting `sdp_session_t` structure.

Return Values: The `sdp_parse()` function returns 0 when the function completes successfully. When the session arguments are invalid, the `sdp_parse()` function returns `EINVAL`. When memory allocation fails while the `sdp_parse()` function is parsing `sdp_info`, the function returns `ENOMEM`. The value of `errno` does not change in the event of an error.

EXAMPLE 2-5 Example: Parsing an SDP Session Structure

In this example, the SDP session structure is as follows:

EXAMPLE 2-5 Example: Parsing an SDP Session Structure (Continued)

```
v=0\r\n
o=jdoe 23423423 234234234 IN IP4 192.168.1.1\r\n
s=SDP seminar\r\n
i=A seminar on the session description protocol\r\n
e=test@host.com
c=IN IP4 156.78.90.1\r\n
t=2873397496 2873404696\r\n
```

After calling the `sdp_parse_t()` function, the resulting `sdp_session_t` structure is as follows:

```
session {
    sdp_session_version = 1
    s_version = 0
    s_origin {
        o_username = "jdoe"
        o_id = 23423423ULL
        o_version = 234234234ULL
        o_nettype = "IN"
        o_addrtype = "IP4"
        o_address = "192.168.1.1"
    }
    s_name = "SDP seminar"
    s_info = "A seminar on the session description protocol"
    s_uri = (nil)
    s_email {
        value = "test@host.com"
        next = (nil)
    }
    s_phone = (nil)
    s_conn {
        c_nettype = "IN"
        c_addrtype = "IP4"
        c_address = "156.78.90.1"
        c_addrcount = 0
        c_ttl = 0
        c_next = (nil)
    }
    s_bw = (nil)
    s_time {
        t_start = 2873397496ULL
        t_stop = 2873404696ULL
        t_repeat = (nil)
        t_next = (nil)
    }
    s_zone = (nil)
    s_key = (nil)
    s_attr = (nil)
    s_media = (nil)
}
```

Cloning an Existing SDP Session Structure

```
sdp_session_t sdp_clone_session(const sdp_session_t *session);
```

The `sdp_clone_session()` function creates a new SDP session structure that is identical to the SDP session structure that is identified in the `session` parameter. The `sdp_clone_session()` function returns the cloned session structure upon successful completion. The `sdp_clone_session()` function returns NULL on failure.

Converting an SDP Session Structure to a String

```
char *sdp_session_to_str(const sdp_session_t *session, int *error);
```

The `sdp_session_to_str()` function returns the string representation of the SDP session structure that is specified by the `session` parameter. The `sdp_session_to_str()` function appends a carriage return/line feed to the end of each SDP field before appending the field to the string.

Return Values: The `sdp_session_to_str()` function returns the string representation of the SDP session structure upon completing successfully. The `sdp_session_to_str()` function returns NULL in all other cases. The `sdp_session_to_str()` function returns an error pointer to EINVAL when the input is null. The `sdp_session_to_str()` function returns an error pointer to ENOMEM when a memory allocation failure occurs. The value of `errno` does not change in the event of an error.

Process Scheduler

This chapter describes the scheduling of processes and how to modify scheduling.

- “[Overview of the Scheduler](#)” on page 43 contains an overview of the scheduler and the time-sharing scheduling class. Other scheduling classes are briefly described.
- “[Commands and Interfaces](#)” on page 47 describes the commands and interfaces that modify scheduling.
- “[Interactions With Other Interfaces](#)” on page 50 describes the effects of scheduling changes on kernel processes and certain interfaces.
- Performance issues to consider when using these commands or interfaces are covered in “[Scheduling and System Performance](#)” on page 51.

The chapter is for developers who need more control over the order of process execution than default scheduling provides. See *Multithreaded Programming Guide* for a description of multithreaded scheduling.

Overview of the Scheduler

When a process is created, the system assigns a lightweight process (LWP) to the process. If the process is multithreaded, more LWPs might be assigned to the process. An LWP is the object that is scheduled by the UNIX system scheduler, which determines when processes run. The scheduler maintains process priorities that are based on configuration parameters, process behavior, and user requests. The scheduler uses these priorities to determine which process runs next. The six priority classes are real-time, system, interactive (IA), fixed-priority (FX), fair-share (FSS), and time-sharing (TS).

The default scheduling is a time-sharing policy. This policy dynamically adjusts process priorities to balance the response time of interactive processes. The policy also dynamically adjusts priorities to balance the throughput of processes that use a lot of CPU time. The time-sharing class has the lowest priority.

The SunOS 5.11 scheduler also provides a real-time scheduling policy. Real-time scheduling enables the assigning of fixed priorities to specific processes by users. The highest-priority real-time user process always gets the CPU as soon as the process is runnable .

The SunOS 5.11 scheduler also provides a policy for fixed-priority scheduling. Fixed-priority scheduling enables the assignment of fixed priorities to specific processes by users. Fixed-priority scheduling uses the same priority range as the time-sharing scheduling class by default.

A program can be written so that its real-time processes have a guaranteed response time from the system. See [Chapter 11, “Real-time Programming and Administration,”](#) for detailed information.

The control of process scheduling provided by real-time scheduling is rarely needed. However, when the requirements for a program include strict timing constraints, real-time processes might be the only way to satisfy those constraints.



Caution – Careless use of real-time processes can have a dramatic negative effect on the performance of time-sharing processes.

Because changes in scheduler administration can affect scheduler behavior, programmers might also need to know something about scheduler administration. The following interfaces affect scheduler administration:

- `dispadmin(1M)` displays or changes scheduler configuration in a running system.
- `ts_dptbl(4)` and `rt_dptbl(4)` are tables that contain the time-sharing and real-time parameters that are used to configure the scheduler.

A process inherits its scheduling parameters, including scheduling class and priority within that class, when the process is created. A process changes class only by user request. The system bases its adjustments of a process' priority on user requests and the policy associated with the scheduler class of the process.

In the default configuration, the initialization process belongs to the time-sharing class. Therefore, all user login shells begin as time-sharing processes.

The scheduler converts class-specific priorities into global priorities. The global priority of a process determines when the process runs. The scheduler always runs the runnable process with the highest global priority. Higher priorities run first. A process assigned to the CPU runs until the process sleeps, uses its time slice, or is preempted by a higher-priority process. Processes with the same priority run in sequence, around a circle.

All real-time processes have higher priorities than any kernel process, and all kernel processes have higher priorities than any time-sharing process.

Note – In a single processor system, no kernel process and no time-sharing process runs while a runnable real-time process exists.

Administrators specify default time slices in the configuration tables. Users can assign per-process time slices to real-time processes.

You can display the global priority of a process with the `-cl` options of the `ps(1)` command. You can display configuration information about class-specific priorities with the `prionctl(1)` command and the `dispadmin(1M)` command.

The following sections describe the scheduling policies of the six scheduling classes.

Time-Sharing Class

The goal of the time-sharing policy is to provide good response time to interactive processes and good throughput to CPU-bound processes. The scheduler switches CPU allocation often enough to provide good response time, but not so often that the system spends too much time on switching. Time slices are typically a few hundred milliseconds.

The time-sharing policy changes priorities dynamically and assigns time slices of different lengths. The scheduler raises the priority of a process that sleeps after only a little CPU use. For example, a process sleeps when the process starts an I/O operation such as a terminal read or a disk read. Frequent sleeps are characteristic of interactive tasks such as editing and running simple shell commands. The time-sharing policy lowers the priority of a process that uses the CPU for long periods without sleeping.

The time-sharing policy that is the default gives larger time slices to processes with lower priorities. A process with a low priority is likely to be CPU-bound. Other processes get the CPU first, but when a low-priority process finally gets the CPU, that process gets a larger time slice. If a higher-priority process becomes runnable during a time slice, however, the higher-priority process preempts the running process.

Global process priorities and user-supplied priorities are in ascending order: higher priorities run first. The user priority runs from the negative of a configuration-dependent maximum to the positive of that maximum. A process inherits its user priority. Zero is the default initial user priority.

The “user priority limit” is the configuration-dependent maximum value of the user priority. You can set a user priority to any value lower than the user priority limit. With appropriate permission, you can raise the user priority limit. Zero is the user priority limit by default.

You can lower the user priority of a process to give the process reduced access to the CPU. Alternately, with the appropriate permission, raise the user priority to get faster service. The

user priority cannot be set to a value that is higher than the user priority limit. Therefore, you must raise the user priority limit before raising the user priority if both have their default values at zero.

An administrator configures the maximum user priority independent of global time-sharing priorities. For example, in the default configuration a user can set a user priority in the -20 to $+20$ range. However, 60 time-sharing global priorities are configured.

The scheduler manages time-sharing processes by using configurable parameters in the time-sharing parameter table `ts_dptbl(4)`. This table contains information specific to the time-sharing class.

System Class

The system class uses a fixed-priority policy to run kernel processes such as servers and housekeeping processes like the paging daemon. The system class is reserved to the kernel. Users cannot add a process to the system class. Users cannot remove a process from the system class. Priorities for system class processes are set up in the kernel code. The priorities of system processes do not change once established. User processes that run in kernel mode are not in the system class.

Real-time Class

The real-time class uses a scheduling policy with fixed priorities so that critical processes run in predetermined order. Real-time priorities never change except when a user requests a change. Privileged users can use the `prionctl(1)` command or the `prionctl(2)` interface to assign real-time priorities.

The scheduler manages real-time processes by using configurable parameters in the real-time parameter table `rt_dptbl(4)`. This table contains information specific to the real-time class.

Interactive Class

The IA class is very similar to the TS class. When used in conjunction with a windowing system, processes have a higher priority while running in a window with the input focus. The IA class is the default class while the system runs a windowing system. The IA class is otherwise identical to the TS class, and the two classes share the same `ts_dptbl` dispatch parameter table.

Fair-Share Class

The FSS class is used by the Fair-Share Scheduler (`FSS(7)`) to manage application performance by explicitly allocating shares of CPU resources to projects. A share indicates a project's entitlement to available CPU resources. The system tracks resource usage over time. The system

reduces entitlement when usage is heavy. The system increases entitlement when usage is light. The FSS schedules CPU time among processes according to their owners' entitlements, independent of the number of processes each project owns. The FSS class uses the same priority range as the TS and IA classes. See the FSS man page for more details.

Fixed-Priority Class

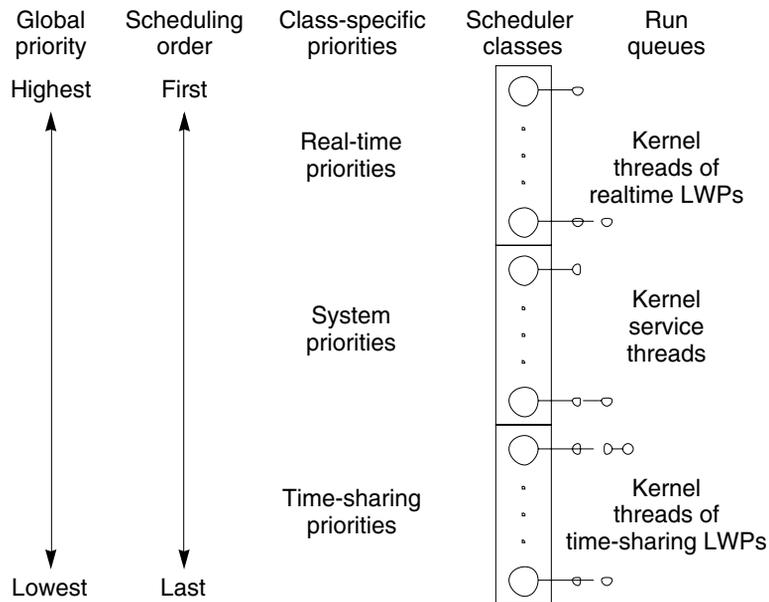
The FX class provides a fixed-priority preemptive scheduling policy. This policy is used by processes that require user or application control of scheduling priorities but are not dynamically adjusted by the system. By default, the FX class has the same priority range as the TS, IA, and FSS classes. The FX class allows user or application control of scheduling priorities through user priority values assigned to processes within the class. These user priority values determine the scheduling priority of a fixed-priority process relative to other processes within its class.

The scheduler manages fixed-priority processes by using configurable parameters in the fixed-priority dispatch parameter table `fx_dptbl(4)`. This table contains information specific to the fixed-priority class.

Commands and Interfaces

The following figure illustrates the default process priorities.

FIGURE 3-1 Process Priorities (Programmer's View)



A process priority has meaning only in the context of a scheduler class. You specify a process priority by specifying a class and a class-specific priority value. The class and class-specific value are mapped by the system into a global priority that the system uses to schedule processes.

A system administrator's view of priorities is different from the view of a user or programmer. When configuring scheduler classes, an administrator deals directly with global priorities. The system maps priorities supplied by users into these global priorities. See *Oracle Solaris Administration: Common Tasks* for more information about priorities.

The `ps(1)` command with `-cel` options reports global priorities for all active processes. The `prctl(1)` command reports the class-specific priorities that users and programmers use.

The `prctl(1)` command and the `prctl(2)` and `prctlset(2)` interfaces are used to set or retrieve scheduler parameters for processes. Setting priorities generally follows the same sequence for the command and both interfaces:

1. Specify the target processes.
2. Specify the scheduler parameters that you want for those processes.
3. Execute the command or interface to set the parameters for the processes.

Process IDs are basic properties of UNIX processes. See `intro(2)` for more information. The class ID is the scheduler class of the process. `prctl(2)` works only for the time-sharing and the real-time classes, not for the system class.

priocntl Usage

The `priocntl(1)` utility performs four different control interfaces on the scheduling of a process:

- `priocntl -l` Displays configuration information
- `priocntl -d` Displays the scheduling parameters of processes
- `priocntl -s` Sets the scheduling parameters of processes
- `priocntl -e` Executes a command with the specified scheduling parameters

The following examples demonstrate the use of `priocntl(1)`.

- The `-l` option for the default configuration produces the following output:

```
$ priocntl -l
CONFIGURED CLASSES
=====

SYS (System Class)

TS (Time Sharing)
Configured TS User Priority Range -60 through 60

RT (Real Time)
Maximum Configured RT Priority: 59
```

- To display information on all processes, do the following:


```
$ priocntl -d -i all
```
- To display information on all time-sharing processes:


```
$ priocntl -d -i class TS
```
- To display information on all processes with user ID 103 or 6626, do the following:


```
$ priocntl -d -i uid 103 6626
```
- To make the process with ID 24668 a real-time process with default parameters, do the following:


```
$ priocntl -s -c RT -i pid 24668
```
- To make 3608 RT with priority 55 and a one-fifth second time slice:


```
$ priocntl -s -c RT -p 55 -t 1 -r 5 -i pid 3608
```
- To change all processes into time-sharing processes, do the following:


```
$ priocntl -s -c TS -i all
```
- To reduce TS user priority and user priority limit to -10 for uid 1122:


```
$ priocntl -s -c TS -p -10 -m -10 -i uid 1122
```
- To start a real-time shell with default real-time priority, do the following:


```
$ priocntl -e -c RT /bin/sh
```

- To run `make` with a time-sharing user priority of `-10`, do the following:

```
$ priocntl -e -c TS -p -10 make bigprog
```

`priocntl(1)` includes the interface of `nice(1)`. `nice` works only on time-sharing processes and uses higher numbers to assign lower priorities. The previous example is equivalent to using `nice(1)` to set an increment of `10`:

```
$ nice -10 make bigprog
```

priocntl Interface

`priocntl(2)` manages the scheduling parameters of a process or set of processes. An invocation of `priocntl(2)` can act on a LWP, on a single process, or on a group of processes. A group of processes can be identified by parent process, process group, session, user, group, class, or all active processes. For more details, see the `priocntl` man page.

The `PC_GETCLINFO` command gets a scheduler class name and parameters when given the class ID. This command enables you to write programs that make no assumptions about what classes are configured.

The `PC_SETXPARMS` command sets the scheduler class and parameters of a set of processes. The `idtype` and `id` input arguments specify the processes to be changed.

Interactions With Other Interfaces

Altering the priority of a process in the `TS` class can affect the behavior of other processes in the `TS` class. This section identifies ways in which a scheduling change can affect other processes.

Kernel Processes

The kernel's daemon and housekeeping processes are members of the system scheduler class. Users can neither add processes to nor remove processes from this class, nor can users change the priorities of these processes. The command `ps -ce1` lists the scheduler class of all processes. A `SYS` entry in the `CLS` column identifies processes in the system class when you run `ps(1)` with the `-f` option.

Using fork and exec

Scheduler class, priority, and other scheduler parameters are inherited across the `fork(2)` and `exec(2)` interfaces.

Using nice

The `nice(1)` command and the `nice(2)` interface work as in previous versions of the UNIX system. These commands enable you to change the priority of a time-sharing process. Use lower numeric values to assign higher time-sharing priorities with these interfaces.

To change the scheduler class of a process or to specify a real-time priority, use `prctl(2)`. Use higher numeric values to assign higher priorities.

init(1M)

The `init(1M)` process is a special case to the scheduler. To change the scheduling properties of `init(1M)`, `init` must be the only process specified by `idtype` and `id` or by the `procset` structure.

Scheduling and System Performance

The scheduler determines when and for how long processes run. Therefore, the scheduler's behavior strongly affects a system's performance.

By default, all user processes are time-sharing processes. A process changes class only by a `prctl(2)` call.

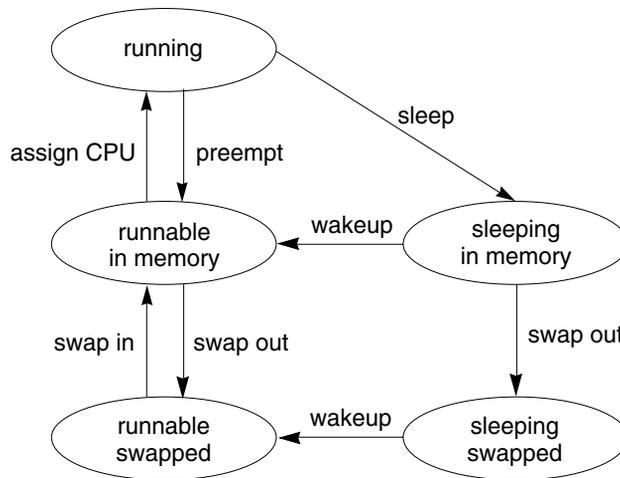
All real-time process priorities have a higher priority than any time-sharing process. Time-sharing processes or system processes cannot run while any real-time process is runnable. A real-time application that occasionally fails to relinquish control of the CPU can completely lock out other users and essential kernel housekeeping.

Besides controlling process class and priorities, a real-time application must also control other factors that affect its performance. The most important factors in performance are CPU power, amount of primary memory, and I/O throughput. These factors interact in complex ways. The `sar(1)` command has options for reporting on all performance factors.

Process State Transition

Applications that have strict real-time constraints might need to prevent processes from being swapped or paged out to secondary memory. A simplified overview of UNIX process states and the transitions between states is shown in the following figure.

FIGURE 3-2 Process State Transition Diagram



An active process is normally in one of the five states in the diagram. The arrows show how the process changes states.

- A process is running if the process is assigned to a CPU. A process is removed from the running state by the scheduler if a process with a higher priority becomes runnable. A process is also preempted if a process of equal priority is runnable when the original process consumes its entire time slice.
 - A process is runnable in memory if the process is in primary memory and ready to run, but is not assigned to a CPU.
 - A process is sleeping in memory if the process is in primary memory but is waiting for a specific event before continuing execution. For example, a process sleeps while waiting for an I/O operation to complete, for a locked resource to be unlocked, or for a timer to expire. When the event occurs, a wakeup call is sent to the process. If the reason for its sleep is gone, the process becomes runnable.
 - When a process' address space has been written to secondary memory, and that process is not waiting for a specific event, the process is runnable and swapped.
 - If a process is waiting for a specific event and has had its whole address space written to secondary memory, the process is sleeping and swapped.
- If a machine does not have enough primary memory to hold all its active processes, that machine must page or swap some address space to secondary memory.
- When the system is short of primary memory, the system writes individual pages of some processes to secondary memory but leaves those processes runnable. When a running process, accesses those pages, the process sleeps while the pages are read back into primary memory.

- When the system encounters a more serious shortage of primary memory, the system writes all the pages of some processes to secondary memory. The system marks the pages that have been written to secondary memory as swapped. Such processes can only be scheduled when the system scheduler daemon selects these processes to be read back into memory.

Both paging and swapping cause delay when a process is ready to run again. For processes that have strict timing requirements, this delay can be unacceptable.

To avoid swapping delays, real-time processes are never swapped, though parts of such processes can be paged. A program can prevent paging and swapping by locking its text and data into primary memory. For more information, see the `memcntl(2)` man page. How much memory can be locked is limited by how much memory is configured. Also, locking too much can cause intolerable delays to processes that do not have their text and data locked into memory.

Trade-offs between the performance of real-time processes and the performance of other processes depend on local needs. On some systems, process locking might be required to guarantee the necessary real-time response.

Note – See “Dispatch Latency” on page 255 for information about latencies in real-time applications.

Locality Group APIs

This chapter describes the APIs that applications use to interact with locality groups.

This chapter discusses the following topics:

- “[Locality Groups Overview](#)” on page 55 describes the locality group abstraction.
- “[Verifying the Interface Version](#)” on page 58 describes the functions that give information about the interface.
- “[Initializing the Locality Group Interface](#)” on page 58 describes function calls that initialize and shut down the portion of the interface that is used to traverse the locality group hierarchy and to discover the contents of a locality group.
- “[Locality Group Hierarchy](#)” on page 60 describes function calls that navigate the locality group hierarchy and functions that get characteristics of the locality group hierarchy.
- “[Locality Group Contents](#)” on page 62 describes function calls that retrieve information about a locality group's contents.
- “[Locality Group Characteristics](#)” on page 64 describes function calls that retrieve information about a locality group's characteristics.
- “[Locality Groups and Thread and Memory Placement](#)” on page 65 describes how to affect the locality group placement of a thread and its memory.
- “[Examples of API Usage](#)” on page 73 contains code that performs example tasks by using the APIs that are described in this chapter.

Locality Groups Overview

Shared memory multiprocessor computers contain multiple CPUs. Each CPU can access all of the memory in the machine. In some shared memory multiprocessors, the memory architecture enables each CPU to access some areas of memory more quickly than other areas.

When a machine with such a memory architecture runs the Oracle Solaris software, providing information to the kernel about the shortest access times between a given CPU and a given area

of memory can improve the system's performance. The locality group (lgroup) abstraction has been introduced to handle this information. The lgroup abstraction is part of the Memory Placement Optimization (MPO) feature.

An lgroup is a set of CPU-like and memory-like devices in which each CPU in the set can access any memory in that set within a bounded latency interval. The value of the latency interval represents the least common latency between all the CPUs and all the memory in that lgroup. The latency bound that defines an lgroup does not restrict the maximum latency between members of that lgroup. The value of the latency bound is the shortest latency that is common to all possible CPU-memory pairs in the group.

Lgroups are hierarchical. The lgroup hierarchy is a Directed Acyclic Graph (DAG) and is similar to a tree, except that an lgroup might have more than one parent. The root lgroup contains all the resources in the system and can include child lgroups. Furthermore, the root lgroup can be characterized as having the highest latency value of all the lgroups in the system. All of its child lgroups will have lower latency values. The lgroups closer to the root have a higher latency while lgroups closer to leaves have lower latency.

A computer in which all the CPUs can access all the memory in the same amount of time can be represented with a single lgroup (see [Figure 4-1](#)). A computer in which some of the CPUs can access some areas of memory in a shorter time than other areas can be represented by using multiple lgroups (see [Figure 4-2](#)).

FIGURE 4-1 Single Locality Group Schematic

Machine with single latency
is represented by one lgroup

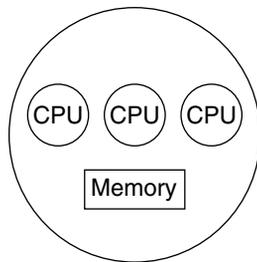
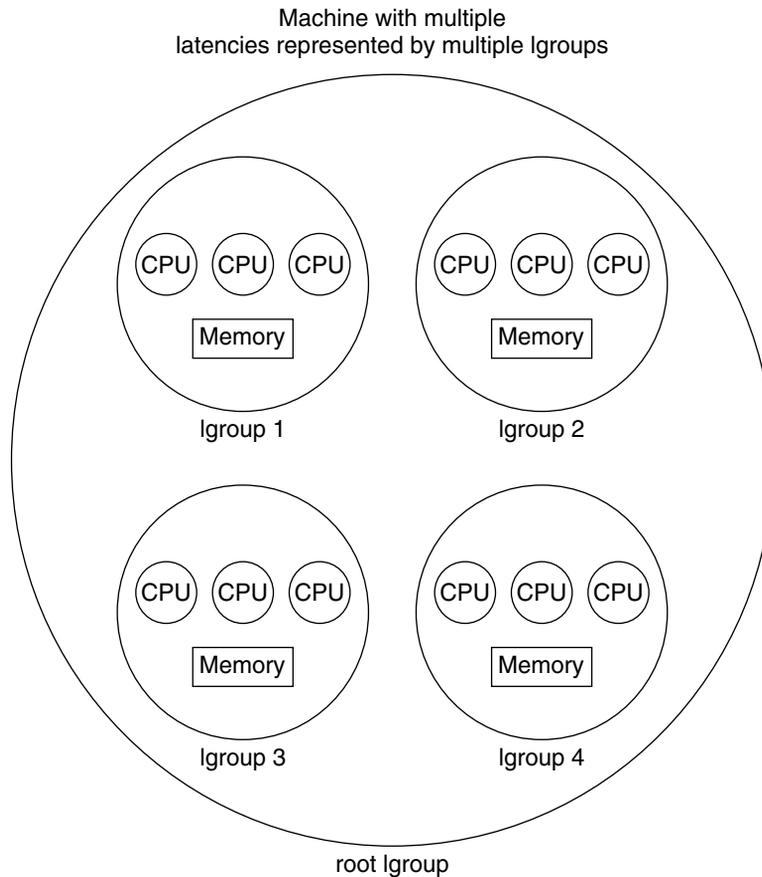


FIGURE 4-2 Multiple Locality Groups Schematic



The organization of the lgroup hierarchy simplifies the task of finding the nearest resources in the system. Each thread is assigned a home lgroup upon creation. The operating system attempts to allocate resources for the thread from the thread's home lgroup by default. For example, the Oracle Solaris kernel attempts to schedule a thread to run on the CPUs in the thread's home lgroup and allocate the thread's memory in the thread's home lgroup by default. If the desired resources are not available from the thread's home lgroup, the kernel can traverse the lgroup hierarchy to find the next nearest resources from parents of the home lgroup. If the desired resources are not available in the home lgroup's parents, the kernel continues to traverse the lgroup hierarchy to the successive ancestor lgroups of the home lgroup. The root lgroup is the ultimate ancestor of all other lgroups in a machine and contains all of the machine's resources.

The lgroup APIs export the lgroup abstraction for applications to use for observability and performance tuning. A new library, called `liblgrp`, contains the new APIs. Applications can use the APIs to perform the following tasks:

- Traverse the group hierarchy
- Discover the contents and characteristics of a given lgroup
- Affect the thread and memory placement on lgroups

Verifying the Interface Version

The `lgrp_version(3LGRP)` function must be used to verify the presence of a supported lgroup interface before using the lgroup API. The `lgrp_version()` function has the following syntax:

```
#include <sys/lgrp_user.h>
int lgrp_version(const int version);
```

The `lgrp_version()` function takes a version number for the lgroup interface as an argument and returns the lgroup interface version that the system supports. When the current implementation of the lgroup API supports the version number in the `version` argument, the `lgrp_version()` function returns that version number. Otherwise, the `lgrp_version()` function returns `LGRP_VER_NONE`.

EXAMPLE 4-1 Example of `lgrp_version()` Use

```
#include <sys/lgrp_user.h>
if (lgrp_version(LGRP_VER_CURRENT) != LGRP_VER_CURRENT) {
    fprintf(stderr, "Built with unsupported lgroup interface %d\n",
            LGRP_VER_CURRENT);
    exit (1);
}
```

Initializing the Locality Group Interface

Applications must call `lgrp_init(3LGRP)` in order to use the APIs for traversing the lgroup hierarchy and to discover the contents of the lgroup hierarchy. The call to `lgrp_init()` gives the application a consistent snapshot of the lgroup hierarchy. The application developer can specify whether the snapshot contains only the resources that are available to the calling thread specifically or the resources that are available to the operating system in general. The `lgrp_init()` function returns a cookie that is used for the following tasks:

- Navigating the lgroup hierarchy
- Determining the contents of an lgroup
- Determining whether the snapshot is current

Using `lgrp_init()`

The `lgrp_init()` function initializes the `lgroup` interface and takes a snapshot of the `lgroup` hierarchy.

```
#include <sys/lgrp_user.h>
lgrp_cookie_t lgrp_init(lgrp_view_t view);
```

When the `lgrp_init()` function is called with `LGRP_VIEW_CALLER` as the view, the function returns a snapshot that contains only the resources that are available to the calling thread. When the `lgrp_init()` function is called with `LGRP_VIEW_OS` as the view, the function returns a snapshot that contains the resources that are available to the operating system. When a thread successfully calls the `lgrp_init()` function, the function returns a cookie that is used by any function that interacts with the `lgroup` hierarchy. When a thread no longer needs the cookie, call the `lgrp_fini()` function with the cookie as the argument.

The `lgroup` hierarchy consists of a root `lgroup` that contains all of the machine's CPU and memory resources. The root `lgroup` might contain other locality groups bounded by smaller latencies.

The `lgrp_init()` function can return two errors. When a view is invalid, the function returns `EINVAL`. When there is insufficient memory to allocate the snapshot of the `lgroup` hierarchy, the function returns `ENOMEM`.

Using `lgrp_fini()`

The `lgrp_fini(3LGRP)` function ends the usage of a given cookie and frees the corresponding `lgroup` hierarchy snapshot.

```
#include <sys/lgrp_user.h>
int lgrp_fini(lgrp_cookie_t cookie);
```

The `lgrp_fini()` function takes a cookie that represents an `lgroup` hierarchy snapshot created by a previous call to `lgrp_init()`. The `lgrp_fini()` function frees the memory that is allocated to that snapshot. After the call to `lgrp_fini()`, the cookie is invalid. Do not use that cookie again.

When the cookie passed to the `lgrp_fini()` function is invalid, `lgrp_fini()` returns `EINVAL`.

Locality Group Hierarchy

The APIs that are described in this section enable the calling thread to navigate the lgroup hierarchy. The lgroup hierarchy is a directed acyclic graph that is similar to a tree, except that a node might have more than one parent. The root lgroup represents the whole machine and contains all of that machine's resources. The root lgroup is the lgroup with the highest latency value in the system. Each of the child lgroups contains a subset of the hardware that is in the root lgroup. Each child lgroup is bounded by a lower latency value. Locality groups that are closer to the root have more resources and a higher latency. Locality groups that are closer to the leaves have fewer resources and a lower latency. An lgroup can contain resources directly within its latency boundary. An lgroup can also contain leaf lgroups that contain their own sets of resources. The resources of leaf lgroups are available to the lgroup that encapsulates those leaf lgroups.

Using `lgrp_cookie_stale()`

The `lgrp_cookie_stale(3LGRP)` function determines whether the snapshot of the lgroup hierarchy represented by the given cookie is current.

```
#include <sys/lgrp_user.h>
int lgrp_cookie_stale(lgrp_cookie_t cookie);
```

The cookie returned by the `lgrp_init()` function can become stale due to several reasons that depend on the view that the snapshot represents. A cookie that is returned by calling the `lgrp_init()` function with the view set to `LGRP_VIEW_OS` can become stale due to changes in the lgroup hierarchy such as dynamic reconfiguration or a change in a CPU's online status. A cookie that is returned by calling the `lgrp_init()` function with the view set to `LGRP_VIEW_CALLER` can become stale due to changes in the calling thread's processor set or changes in the lgroup hierarchy. A stale cookie is refreshed by calling the `lgrp_fini()` function with the old cookie, followed by calling `lgrp_init()` to generate a new cookie.

The `lgrp_cookie_stale()` function returns `EINVAL` when the given cookie is invalid.

Using `lgrp_view()`

The `lgrp_view(3LGRP)` function determines the view with which a given lgroup hierarchy snapshot was taken.

```
#include <sys/lgrp_user.h>
lgrp_view_t lgrp_view(lgrp_cookie_t cookie);
```

The `lgrp_view()` function takes a cookie that represents a snapshot of the lgroup hierarchy and returns the snapshot's view of the lgroup hierarchy. Snapshots that are taken with the view

LGRP_VIEW_CALLER contain only the resources that are available to the calling thread. Snapshots that are taken with the view LGRP_VIEW_OS contain all the resources that are available to the operating system.

The `lgrp_view()` function returns `EINVAL` when the given cookie is invalid.

Using `lgrp_nlgrps()`

The `lgrp_nlgrps(3LGRP)` function returns the number of locality groups in the system. If a system has only one locality group, memory placement optimizations have no effect.

```
#include <sys/lgrp_user.h>
int lgrp_nlgrps(lgrp_cookie_t cookie);
```

The `lgrp_nlgrps()` function takes a cookie that represents a snapshot of the lgroup hierarchy and returns the number of lgroups available in the hierarchy.

The `lgrp_nlgrps()` function returns `EINVAL` when the cookie is invalid.

Using `lgrp_root()`

The `lgrp_root(3LGRP)` function returns the root lgroup ID.

```
#include <sys/lgrp_user.h>
lgrp_id_t lgrp_root(lgrp_cookie_t cookie);
```

The `lgrp_root()` function takes a cookie that represents a snapshot of the lgroup hierarchy and returns the root lgroup ID.

Using `lgrp_parents()`

The `lgrp_parents(3LGRP)` function takes a cookie that represents a snapshot of the lgroup hierarchy and returns the number of parent lgroups for the specified lgroup.

```
#include <sys/lgrp_user.h>
int lgrp_parents(lgrp_cookie_t cookie, lgrp_id_t child,
               lgrp_id_t *lgrp_array, uint_t lgrp_array_size);
```

If `lgrp_array` is not `NULL` and the value of `lgrp_array_size` is not zero, the `lgrp_parents()` function fills the array with parent lgroup IDs until the array is full or all parent lgroup IDs are in the array. The root lgroup has zero parents. When the `lgrp_parents()` function is called for the root lgroup, `lgrp_array` is not filled in.

The `lgrp_parents()` function returns `EINVAL` when the cookie is invalid. The `lgrp_parents()` function returns `ESRCH` when the specified lgroup ID is not found.

Using `lgrp_children()`

The `lgrp_children(3LGRP)` function takes a cookie that represents the calling thread's snapshot of the lgroup hierarchy and returns the number of child lgroups for the specified lgroup.

```
#include <sys/lgrp_user.h>
int lgrp_children(lgrp_cookie_t cookie, lgrp_id_t parent,
                 lgrp_id_t *lgrp_array, uint_t lgrp_array_size);
```

If `lgrp_array` is not NULL and the value of `lgrp_array_size` is not zero, the `lgrp_children()` function fills the array with child lgroup IDs until the array is full or all child lgroup IDs are in the array.

The `lgrp_children()` function returns `EINVAL` when the cookie is invalid. The `lgrp_children()` function returns `ESRCH` when the specified lgroup ID is not found.

Locality Group Contents

The following APIs retrieve information about the contents of a given lgroup.

The lgroup hierarchy organizes the domain's resources to simplify the process of locating the nearest resource. Leaf lgroups are defined with resources that have the least latency. Each of the successive ancestor lgroups of a given leaf lgroup contains the next nearest resources to its child lgroup. The root lgroup contains all of the resources that are in the domain.

The resources of a given lgroup are contained directly within that lgroup or indirectly within the leaf lgroups that the given lgroup encapsulates. Leaf lgroups directly contain their resources and do not encapsulate any other lgroups.

Using `lgrp_resources()`

The `lgrp_resources()` function returns the number of resources contained in a specified lgroup.

```
#include <sys/lgrp_user.h>
int lgrp_resources(lgrp_cookie_t cookie, lgrp_id_t lgrp, lgrp_id_t *lgrpid,
                  uint_t count, lgrp_rsrc_t type);
```

The `lgrp_resources()` function takes a cookie that represents a snapshot of the lgroup hierarchy. That cookie is obtained from the `lgrp_init()` function. The `lgrp_resources()` function returns the number of resources that are in the lgroup with the ID that is specified by the value of the `lgrp` argument. The `lgrp_resources()` function represents the resources with a set of lgroups that directly contain CPU or memory resources. The `lgrp_rsrc_t` argument can have the following two values:

`LGRP_RSRC_CPU` The `lgrp_resources()` function returns the number of CPU resources.

`LGRP_RSRC_MEM` The `lgrp_resources()` function returns the number of memory resources.

When the value passed in the `lgrpids[]` argument is not null and the `count` argument is not zero, the `lgrp_resources()` function stores lgroup IDs in the `lgrpids[]` array. The number of lgroup IDs stored in the array can be up to the value of the `count` argument.

The `lgrp_resources()` function returns `EINVAL` when the specified cookie, lgroup ID, or type are not valid. The `lgrp_resources()` function returns `ESRCH` when the function does not find the specified lgroup ID.

Using `lgrp_cpus()`

The `lgrp_cpus(3LGRP)` function takes a cookie that represents a snapshot of the lgroup hierarchy and returns the number of CPUs in a given lgroup.

```
#include <sys/lgrp_user.h>
int lgrp_cpus(lgrp_cookie_t cookie, lgrp_id_t lgrp, processorid_t *cpuids,
             uint_t count, int content);
```

If the `cpuid[]` argument is not `NULL` and the CPU count is not zero, the `lgrp_cpus()` function fills the array with CPU IDs until the array is full or all the CPU IDs are in the array.

The `content` argument can have the following two values:

`LGRP_CONTENT_ALL` The `lgrp_cpus()` function returns IDs for the CPUs in this lgroup and this lgroup's descendants.

`LGRP_CONTENT_DIRECT` The `lgrp_cpus()` function returns IDs for the CPUs in this lgroup only.

The `lgrp_cpus()` function returns `EINVAL` when the cookie, lgroup ID, or one of the flags is not valid. The `lgrp_cpus()` function returns `ESRCH` when the specified lgroup ID is not found.

Using `lgrp_mem_size()`

The `lgrp_mem_size(3LGRP)` function takes a cookie that represents a snapshot of the lgroup hierarchy and returns the size of installed or free memory in the given lgroup. The `lgrp_mem_size()` function reports memory sizes in bytes.

```
#include <sys/lgrp_user.h>
lgrp_mem_size_t lgrp_mem_size(lgrp_cookie_t cookie, lgrp_id_t lgrp,
                             int type, int content)
```

The `type` argument can have the following two values:

LGRP_MEM_SZ_FREE	The <code>lgrp_mem_size()</code> function returns the amount of free memory in bytes.
LGRP_MEM_SZ_INSTALLED	The <code>lgrp_mem_size()</code> function returns the amount of installed memory in bytes.

The *content* argument can have the following two values:

LGRP_CONTENT_ALL	The <code>lgrp_mem_size()</code> function returns the amount of memory in this lgroup and this lgroup's descendants.
LGRP_CONTENT_DIRECT	The <code>lgrp_mem_size()</code> function returns the amount of memory in this lgroup only.

The `lgrp_mem_size()` function returns `EINVAL` when the cookie, lgroup ID, or one of the flags is not valid. The `lgrp_mem_size()` function returns `ESRCH` when the specified lgroup ID is not found.

Locality Group Characteristics

The following API retrieves information about the characteristics of a given lgroup.

Using `lgrp_latency_cookie()`

The `lgrp_latency(3LGRP)` function returns the latency between a CPU in one lgroup to the memory in another lgroup.

```
#include <sys/lgrp_user.h>
int lgrp_latency_cookie(lgrp_cookie_t cookie, lgrp_id_t from, lgrp_id_t to,
                       lat_between_t between);
```

The `lgrp_latency_cookie()` function takes a cookie that represents a snapshot of the lgroup hierarchy. The `lgrp_init()` function creates this cookie. The `lgrp_latency_cookie()` function returns a value that represents the latency between a hardware resource in the lgroup given by the value of the *from* argument and a hardware resource in the lgroup given by the value of the *to* argument. If both arguments point to the same lgroup, the `lgrp_latency_cookie()` function returns the latency value within that lgroup.

Note – The latency value returned by the `lgrp_latency_cookie()` function is defined by the operating system and is platform-specific. This value does not necessarily represent the actual latency between hardware devices. Use this value only for comparison within one domain.

When the value of the *between* argument is `LGRP_LAT_CPU_TO_MEM`, the `lgrp_latency_cookie()` function measures the latency from a CPU resource to a memory resource.

The `lgrp_latency_cookie()` function returns `EINVAL` when the lgroup ID is not valid. When the `lgrp_latency_cookie()` function does not find the specified lgroup ID, the “from” lgroup does not contain any CPUs, or the “to” lgroup does not have any memory, the `lgrp_latency_cookie()` function returns `ESRCH`.

Locality Groups and Thread and Memory Placement

This section discusses the APIs used to discover and affect thread and memory placement with respect to lgroups.

- The `lgrp_home(3LGRP)` function is used to discover thread placement.
- The `meminfo(2)` system call is used to discover memory placement.
- The `MADV_ACCESS` flags to the `madvise(3C)` function are used to affect memory allocation among lgroups.
- The `lgrp_affinity_set(3LGRP)` function can affect thread and memory placement by setting a thread's affinity for a given lgroup.
- The affinities of an lgroup may specify an order of preference for lgroups from which to allocate resources.
- The kernel needs information about the likely pattern of an application's memory use in order to allocate memory resources efficiently.
- The `madvise()` function and its shared object analogue `madv.so.1` provide this information to the kernel.
- A running process can gather memory usage information about itself by using the `meminfo()` system call.

Using `lgrp_home()`

The `lgrp_home()` function returns the home lgroup for the specified process or thread.

```
#include <sys/lgrp_user.h>
lgrp_id_t lgrp_home(idtype_t idtype, id_t id);
```

The `lgrp_home()` function returns `EINVAL` when the ID type is not valid. The `lgrp_home()` function returns `EPERM` when the effective user of the calling process is not the superuser and the

real or effective user ID of the calling process does not match the real or effective user ID of one of the threads. The `lgrp_home()` function returns `ESRCH` when the specified process or thread is not found.

Using `madvise()`

The `madvise()` function advises the kernel that a region of user virtual memory in the range starting at the address specified in `addr` and with length equal to the value of the `len` parameter is expected to follow a particular pattern of use. The kernel uses this information to optimize the procedure for manipulating and maintaining the resources associated with the specified range. Use of the `madvise()` function can increase system performance when used by programs that have specific knowledge of their access patterns over memory.

```
#include <sys/types.h>
#include <sys/mman.h>
int madvise(caddr_t addr, size_t len, int advice);
```

The `madvise()` function provides the following flags to affect how a thread's memory is allocated among lgroups:

<code>MADV_ACCESS_DEFAULT</code>	This flag resets the kernel's expected access pattern for the specified range to the default.
<code>MADV_ACCESS_LWP</code>	This flag advises the kernel that the next LWP to touch the specified address range is the LWP that will access that range the most. The kernel allocates the memory and other resources for this range and the LWP accordingly.
<code>MADV_ACCESS_MANY</code>	This flag advises the kernel that many processes or LWPs will access the specified address range randomly across the system. The kernel allocates the memory and other resources for this range accordingly.

The `madvise()` function can return the following values:

<code>EAGAIN</code>	Some or all of the mappings in the specified address range, from <code>addr</code> to <code>addr+len</code> , are locked for I/O.
<code>EINVAL</code>	The value of the <code>addr</code> parameter is not a multiple of the page size as returned by <code>sysconf(3C)</code> , the length of the specified address range is less than or equal to zero, or the advice is invalid.
<code>EIO</code>	An I/O error occurs while reading from or writing to the file system.
<code>ENOMEM</code>	Addresses in the specified address range are outside the valid range for the address space of a process or the addresses in the specified address range specify one or more pages that are not mapped.

ESTALE The NFS file handle is stale.

Using madv.so.1

The `madv.so.1` shared object enables the selective configuration of virtual memory advice for launched processes and their descendants. To use the shared object, the following string must be present in the environment:

```
LD_PRELOAD=$LD_PRELOAD:madv.so.1
```

The `madv.so.1` shared object applies memory advice as specified by the value of the `MADV` environment variable. The `MADV` environment variable specifies the virtual memory advice to use for all heap, shared memory, and `mmap` regions in the process address space. This advice is applied to all created processes. The following values of the `MADV` environment variable affect resource allocation among lgroups:

<code>access_default</code>	This value resets the kernel's expected access pattern to the default.
<code>access_lwp</code>	This value advises the kernel that the next LWP to touch an address range is the LWP that will access that range the most. The kernel allocates the memory and other resources for this range and the LWP accordingly.
<code>access_many</code>	This value advises the kernel that many processes or LWPs will access memory randomly across the system. The kernel allocates the memory and other resources accordingly.

The value of the `MADVCFGFILE` environment variable is the name of a text file that contains one or more memory advice configuration entries in the form *exec-name:advice-opts*.

The value of *exec-name* is the name of an application or executable. The value of *exec-name* can be a full pathname, a base name, or a pattern string.

The value of *advice-opts* is of the form *region=advice*. The values of *advice* are the same as the values for the `MADV` environment variable. Replace *region* with any of the following legal values:

<code>madv</code>	Advice applies to all heap, shared memory, and <code>mmap(2)</code> regions in the process address space.
<code>heap</code>	The heap is defined to be the <code>brk(2)</code> area. Advice applies to the existing heap and to any additional heap memory allocated in the future.
<code>shm</code>	Advice applies to shared memory segments. See <code>shmat(2)</code> for more information on shared memory operations.
<code>ism</code>	Advice applies to shared memory segments that are using the <code>SHM_SHARE_MMU</code> flag. The <code>ism</code> option takes precedence over <code>shm</code> .

dsm	Advice applies to shared memory segments that are using the SHM_PAGEABLE flag. The dsm option takes precedence over shm.
mapshared	Advice applies to mappings established by the mmap () system call using the MAP_SHARED flag.
mapprivate	Advice applies to mappings established by the mmap () system call using the MAP_PRIVATE flag.
mapanon	Advice applies to mappings established by the mmap () system call using the MAP_ANON flag. The mapanon option takes precedence when multiple options apply.

The value of the MADVERRFILE environment variable is the name of the path where error messages are logged. In the absence of a MADVERRFILE location, the madv . so . 1 shared object logs errors by using `syslog(3C)` with a LOG_ERR as the severity level and LOG_USER as the facility descriptor.

Memory advice is inherited. A child process has the same advice as its parent. The advice is set back to the system default advice after a call to `exec(2)` unless a different level of advice is configured using the madv . so . 1 shared object. Advice is only applied to mmap () regions explicitly created by the user program. Regions established by the run-time linker or by system libraries that make direct system calls are not affected.

madv.so.1 Usage Examples

The following examples illustrate specific aspects of the madv . so . 1 shared object.

EXAMPLE 4-2 Setting Advice for a Set of Applications

This configuration applies advice to all ISM segments for applications with exec names that begin with foo.

```
$ LD_PRELOAD=$LD_PRELOAD:madv.so.1
$ MADVCFGFILE=madvcfg
$ export LD_PRELOAD MADVCFGFILE
$ cat $MADVCFGFILE
foo*:ism=access_lwp
```

EXAMPLE 4-3 Excluding a Set of Applications From Advice

This configuration sets advice for all applications with the exception of ls.

```
$ LD_PRELOAD=$LD_PRELOAD:madv.so.1
$ MADV=access_many
$ MADVCFGFILE=madvcfg
$ export LD_PRELOAD MADV MADVCFGFILE
$ cat $MADVCFGFILE
ls:
```

EXAMPLE 4-4 Pattern Matching in a Configuration File

Because the configuration specified in `MADVCFGFILE` takes precedence over the value set in `MADV`, specifying `*` as the *exec-name* of the last configuration entry is equivalent to setting `MADV`. This example is equivalent to the previous example.

```
$ LD_PRELOAD=$LD_PRELOAD:madv.so.1
$ MADVCFGFILE=madvcfg
$ export LD_PRELOAD MADVCFGFILE
$ cat $MADVCFGFILE
  ls:
  *:madv=access_many
```

EXAMPLE 4-5 Advice for Multiple Regions

This configuration applies one type of advice for `mmap()` regions and different advice for heap and shared memory regions for applications whose `exec()` names begin with `foo`.

```
$ LD_PRELOAD=$LD_PRELOAD:madv.so.1
$ MADVCFGFILE=madvcfg
$ export LD_PRELOAD MADVCFGFILE
$ cat $MADVCFGFILE
  foo*:madv=access_many,heap=sequential,shm=access_lwp
```

Using `meminfo()`

The `meminfo()` function gives the calling process information about the virtual memory and physical memory that the system has allocated to that process.

```
#include <sys/types.h>
#include <sys/mman.h>
int meminfo(const uint64_t inaddr[], int addr_count,
            const uint_t info_req[], int info_count, uint64_t outdata[],
            uint_t validity[]);
```

The `meminfo()` function can return the following types of information:

<code>MEMINFO_VPHYSICAL</code>	The physical memory address corresponding to the given virtual address
<code>MEMINFO_VLGRP</code>	The lgroup to which the physical page corresponding to the given virtual address belongs
<code>MEMINFO_VPAGESIZE</code>	The size of the physical page corresponding to the given virtual address
<code>MEMINFO_VREPLCNT</code>	The number of replicated physical pages that correspond to the given virtual address
<code>MEMINFO_VREPL n</code>	The <i>n</i> th physical replica of the given virtual address

`MEMINFO_VREPL_LGRP | n` The lgroup to which the *n*th physical replica of the given virtual address belongs

`MEMINFO_PLGRP` The lgroup to which the given physical address belongs

The `meminfo()` function takes the following parameters:

inaddr An array of input addresses.

addr_count The number of addresses that are passed to `meminfo()`.

info_req An array that lists the types of information that are being requested.

info_count The number of pieces of information that are requested for each address in the *inaddr* array.

outdata An array where the `meminfo()` function places the results. The array's size is equal to the product of the values of the *info_req* and *addr_count* parameters.

validity An array of size equal to the value of the *addr_count* parameter. The *validity* array contains bitwise result codes. The 0th bit of the result code evaluates the validity of the corresponding input address. Each successive bit in the result code evaluates the validity of the response to the members of the *info_req* array in turn.

The `meminfo()` function returns `EFAULT` when the area of memory to which the *outdata* or *validity* arrays point cannot be written to. The `meminfo()` function returns `EFAULT` when the area of memory to which the *info_req* or *inaddr* arrays point cannot be read from. The `meminfo()` function returns `EINVAL` when the value of *info_count* exceeds 31 or is less than 1. The `meminfo()` function returns `EINVAL` when the value of *addr_count* is less than zero.

EXAMPLE 4-6 Use of `meminfo()` to Print Out Physical Pages and Page Sizes Corresponding to a Set of Virtual Addresses

```
void
print_info(void **addrvec, int how_many)
{
    static const int info[] = {
        MEMINFO_VPHYSICAL,
        MEMINFO_VPAGESIZE};
    uint64_t * inaddr = alloca(sizeof(uint64_t) * how_many);
    uint64_t * outdata = alloca(sizeof(uint64_t) * how_many * 2);
    uint_t * validity = alloca(sizeof(uint_t) * how_many);

    int i;

    for (i = 0; i < how_many; i++)
        inaddr[i] = (uint64_t *)addr[i];

    if (meminfo(inaddr, how_many, info,
               sizeof (info) / sizeof (info[0]),
               outdata, validity) < 0)
        ...
}
```

EXAMPLE 4-6 Use of `meminfo()` to Print Out Physical Pages and Page Sizes Corresponding to a Set of Virtual Addresses (Continued)

```

for (i = 0; i < how_many; i++) {
    if (validity[i] & 1 == 0)
        printf("address 0x%llx not part of address
                space\n",
                inaddr[i]);

    else if (validity[i] & 2 == 0)
        printf("address 0x%llx has no physical page
                associated with it\n",
                inaddr[i]);

    else {
        char buff[80];
        if (validity[i] & 4 == 0)
            strcpy(buff, "<Unknown>", sizeof(buff));
        else
            sprintf(buff, sizeof(buff), "%lld", outdata[i * 2 +
                1]);
        printf("address 0x%llx is backed by physical
                page 0x%llx of size %s\n",
                inaddr[i], outdata[i * 2], buff);
    }
}
}

```

Locality Group Affinity

The kernel assigns a thread to a locality group when the lightweight process (LWP) for that thread is created. That lgroup is called the thread's *home lgroup*. The kernel runs the thread on the CPUs in the thread's home lgroup and allocates memory from that lgroup whenever possible. If resources from the home lgroup are unavailable, the kernel allocates resources from other lgroups. When a thread has affinity for more than one lgroup, the operating system allocates resources from lgroups chosen in order of affinity strength. Lgroups can have one of three distinct affinity levels:

1. `LGRP_AFF_STRONG` indicates strong affinity. If this lgroup is the thread's home lgroup, the operating system avoids rehomeing the thread to another lgroup if possible. Events such as dynamic reconfiguration, processor, offlining, processor binding, and processor set binding and manipulation might still result in thread rehomeing.
2. `LGRP_AFF_WEAK` indicates weak affinity. If this lgroup is the thread's home lgroup, the operating system rehomees the thread if necessary for load balancing purposes.
3. `LGRP_AFF_NONE` indicates no affinity. If a thread has no affinity to any lgroup, the operating system assigns a home lgroup to the thread.

The operating system uses lgroup affinities as advice when allocating resources for a given thread. The advice is factored in with the other system constraints. Processor binding and processor sets do not change lgroup affinities, but might restrict the lgroups on which a thread can run.

Using `lgrp_affinity_get()`

The `lgrp_affinity_get(3LGRP)` function returns the affinity that a LWP has for a given lgroup.

```
#include <sys/lgrp_user.h>
lgrp_affinity_t lgrp_affinity_get(idtype_t idtype, id_t id, lgrp_id_t lgrp);
```

The *idtype* and *id* arguments specify the LWP that the `lgrp_affinity_get()` function examines. If the value of *idtype* is `P_PID`, the `lgrp_affinity_get()` function gets the lgroup affinity for one of the LWPs in the process whose process ID matches the value of the *id* argument. If the value of *idtype* is `P_LWPID`, the `lgrp_affinity_get()` function gets the lgroup affinity for the LWP of the current process whose LWP ID matches the value of the *id* argument. If the value of *idtype* is `P_MYID`, the `lgrp_affinity_get()` function gets the lgroup affinity for the current LWP.

The `lgrp_affinity_get()` function returns `EINVAL` when the given lgroup or ID type is not valid. The `lgrp_affinity_get()` function returns `EPERM` when the effective user of the calling process is not the superuser and the ID of the calling process does not match the real or effective user ID of one of the LWPs. The `lgrp_affinity_get()` function returns `ESRCH` when a given lgroup or LWP is not found.

Using `lgrp_affinity_set()`

The `lgrp_affinity_set(3LGRP)` function sets the affinity that a LWP or set of LWPs have for a given lgroup.

```
#include <sys/lgrp_user.h>
int lgrp_affinity_set(idtype_t idtype, id_t id, lgrp_id_t lgrp,
                    lgrp_affinity_t affinity);
```

The *idtype* and *id* arguments specify the LWP or set of LWPs the `lgrp_affinity_set()` function examines. If the value of *idtype* is `P_PID`, the `lgrp_affinity_set()` function sets the lgroup affinity for all of the LWPs in the process whose process ID matches the value of the *id* argument to the affinity level specified in the *affinity* argument. If the value of *idtype* is `P_LWPID`, the `lgrp_affinity_set()` function sets the lgroup affinity for the LWP of the current process whose LWP ID matches the value of the *id* argument to the affinity level specified in the *affinity* argument. If the value of *idtype* is `P_MYID`, the `lgrp_affinity_set()` function sets the lgroup affinity for the current LWP or process to the affinity level specified in the *affinity* argument.

The `lgrp_affinity_set()` function returns `EINVAL` when the given lgroup, affinity, or ID type is not valid. The `lgrp_affinity_set()` function returns `EPERM` when the effective user of the calling process is not the superuser and the ID of the calling process does not match the real or effective user ID of one of the LWPs. The `lgrp_affinity_set()` function returns `ESRCH` when a given lgroup or LWP is not found.

Examples of API Usage

This section contains code for example tasks that use the APIs that are described in this chapter.

EXAMPLE 4-7 Move Memory to a Thread

The following code sample moves the memory in the address range between *addr* and *addr+len* near the next thread to touch that range.

```
#include <stdio.h>
#include <sys/mman.h>
#include <sys/types.h>

/*
 * Move memory to thread
 */
void
mem_to_thread(caddr_t addr, size_t len)
{
    if (madvise(addr, len, MADV_ACCESS_LWP) < 0)
        perror("madvise");
}
```

EXAMPLE 4-8 Move a Thread to Memory

This sample code uses the `meminfo()` function to determine the lgroup of the physical memory backing the virtual page at the given address. The sample code then sets a strong affinity for that lgroup in an attempt to move the current thread near that memory.

```
#include <stdio.h>
#include <sys/lgrp_user.h>
#include <sys/mman.h>
#include <sys/types.h>

/*
 * Move a thread to memory
 */
int
thread_to_memory(caddr_t va)
{
    uint64_t    addr;
    ulong_t    count;
    lgrp_id_t   home;
    uint64_t    lgrp;
```

EXAMPLE 4-8 Move a Thread to Memory (Continued)

```

uint_t    request;
uint_t    valid;

addr = (uint64_t)va;
count = 1;
request = MEMINFO_VLGRP;
if (meminfo(&addr, 1, &request, 1, &lgrp, &valid) != 0) {
    perror("meminfo");
    return (1);
}

if (lgrp_affinity_set(P_LWPID, P_MYID, lgrp, LGRP_AFF_STRONG) != 0) {
    perror("lgrp_affinity_set");
    return (2);
}

home = lgrp_home(P_LWPID, P_MYID);
if (home == -1) {
    perror ("lgrp_home");
    return (3);
}

if (home != lgrp)
    return (-1);

return (0);
}

```

EXAMPLE 4-9 Walk the lgroup Hierarchy

The following sample code walks through and prints out the lgroup hierarchy.

```

#include <stdio.h>
#include <stdlib.h>
#include <sys/lgrp_user.h>
#include <sys/types.h>

/*
 * Walk and print lgroup hierarchy from given lgroup
 * through all its descendants
 */
int
lgrp_walk(lgrp_cookie_t cookie, lgrp_id_t lgrp, lgrp_content_t content)
{
    lgrp_affinity_t    aff;
    lgrp_id_t          *children;
    processorid_t      *cpuids;
    int                i;
    int                ncpus;
    int                nchildren;
    int                nparents;
    lgrp_id_t          *parents;
    lgrp_mem_size_t    size;

```

EXAMPLE 4-9 Walk the lgroup Hierarchy (Continued)

```

/*
 * Print given lgroup, caller's affinity for lgroup,
 * and desired content specified
 */
printf("LGROUP #d:\n", lgrp);

aff = lgrp_affinity_get(P_LWPID, P_MYID, lgrp);
if (aff == -1)
    perror("lgrp_affinity_get");
printf("\tAFFINITY: %d\n", aff);

printf("CONTENT %d:\n", content);

/*
 * Get CPUs
 */
ncpus = lgrp_cpus(cookie, lgrp, NULL, 0, content);
printf("\t%d CPUS: ", ncpus);
if (ncpus == -1) {
    perror("lgrp_cpus");
    return (-1);
} else if (ncpus > 0) {
    cpuids = malloc(ncpus * sizeof (processorid_t));
    ncpus = lgrp_cpus(cookie, lgrp, cpuids, ncpus, content);
    if (ncpus == -1) {
        free(cpuids);
        perror("lgrp_cpus");
        return (-1);
    }
    for (i = 0; i < ncpus; i++)
        printf("%d ", cpuids[i]);
    free(cpuids);
}
printf("\n");

/*
 * Get memory size
 */
printf("\tMEMORY: ");
size = lgrp_mem_size(cookie, lgrp, LGRP_MEM_SZ_INSTALLED, content);
if (size == -1) {
    perror("lgrp_mem_size");
    return (-1);
}
printf("installed bytes 0x%llx, ", size);
size = lgrp_mem_size(cookie, lgrp, LGRP_MEM_SZ_FREE, content);
if (size == -1) {
    perror("lgrp_mem_size");
    return (-1);
}
printf("free bytes 0x%llx\n", size);

/*
 * Get parents
 */

```

EXAMPLE 4-9 Walk the lgroup Hierarchy (Continued)

```

nparents = lgrp_parents(cookie, lgrp, NULL, 0);
printf("\t%d PARENTS: ", nparents);
if (nparents == -1) {
    perror("lgrp_parents");
    return (-1);
} else if (nparents > 0) {
    parents = malloc(nparents * sizeof (lgrp_id_t));
    nparents = lgrp_parents(cookie, lgrp, parents, nparents);
    if (nparents == -1) {
        free(parents);
        perror("lgrp_parents");
        return (-1);
    }
    for (i = 0; i < nparents; i++)
        printf("%d ", parents[i]);
    free(parents);
}
printf("\n");

/*
 * Get children
 */
nchildren = lgrp_children(cookie, lgrp, NULL, 0);
printf("\t%d CHILDREN: ", nchildren);
if (nchildren == -1) {
    perror("lgrp_children");
    return (-1);
} else if (nchildren > 0) {
    children = malloc(nchildren * sizeof (lgrp_id_t));
    nchildren = lgrp_children(cookie, lgrp, children, nchildren);
    if (nchildren == -1) {
        free(children);
        perror("lgrp_children");
        return (-1);
    }
    printf("Children: ");
    for (i = 0; i < nchildren; i++)
        printf("%d ", children[i]);
    printf("\n");

    for (i = 0; i < nchildren; i++)
        lgrp_walk(cookie, children[i], content);

    free(children);
}
printf("\n");

return (0);
}

```

EXAMPLE 4-10 Find the Closest lgroup With Available Memory Outside a Given lgroup

```

#include <stdio.h>
#include <stdlib.h>
#include <sys/lgrp_user.h>

```

EXAMPLE 4-10 Find the Closest lgroup With Available Memory Outside a Given lgroup (Continued)

```

#include <sys/types.h>

#define    INT_MAX    2147483647

/*
 * Find next closest lgroup outside given one with available memory
 */
lgrp_id_t
lgrp_next_nearest(lgrp_cookie_t cookie, lgrp_id_t from)
{
    lgrp_id_t    closest;
    int          i;
    int          latency;
    int          lowest;
    int          nparents;
    lgrp_id_t    *parents;
    lgrp_mem_size_t    size;

    /*
     * Get number of parents
     */
    nparents = lgrp_parents(cookie, from, NULL, 0);
    if (nparents == -1) {
        perror("lgrp_parents");
        return (LGRP_NONE);
    }

    /*
     * No parents, so current lgroup is next nearest
     */
    if (nparents == 0) {
        return (from);
    }

    /*
     * Get parents
     */
    parents = malloc(nparents * sizeof (lgrp_id_t));
    nparents = lgrp_parents(cookie, from, parents, nparents);
    if (nparents == -1) {
        perror("lgrp_parents");
        free(parents);
        return (LGRP_NONE);
    }

    /*
     * Find closest parent (ie. the one with lowest latency)
     */
    closest = LGRP_NONE;
    lowest = INT_MAX;
    for (i = 0; i < nparents; i++) {
        lgrp_id_t    lgrp;

        /*
         * See whether parent has any free memory

```

EXAMPLE 4-10 Find the Closest lgroup With Available Memory Outside a Given lgroup (Continued)

```

    */
    size = lgrp_mem_size(cookie, parents[i], LGRP_MEM_SZ_FREE,
        LGRP_CONTENT_ALL);
    if (size > 0)
        lgrp = parents[i];
    else {
        if (size == -1)
            perror("lgrp_mem_size");

        /*
         * Find nearest ancestor if parent doesn't
         * have any memory
         */
        lgrp = lgrp_next_nearest(cookie, parents[i]);
        if (lgrp == LGRP_NONE)
            continue;
    }

    /*
     * Get latency within parent lgroup
     */
    latency = lgrp_latency_cookie(lgrp, lgrp);
    if (latency == -1) {
        perror("lgrp_latency_cookie");
        continue;
    }

    /*
     * Remember lgroup with lowest latency
     */
    if (latency < lowest) {
        closest = lgrp;
        lowest = latency;
    }
}

free(parents);
return (closest);
}

/*
 * Find lgroup with memory nearest home lgroup of current thread
 */
lgrp_id_t
lgrp_nearest(lgrp_cookie_t cookie)
{
    lgrp_id_t    home;
    longlong_t  size;

    /*
     * Get home lgroup
     */
    home = lgrp_home(P_LWPID, P_MYID);

    /*

```

EXAMPLE 4-10 Find the Closest lgroup With Available Memory Outside a Given lgroup *(Continued)*

```

    * See whether home lgroup has any memory available in its hierarchy
    */
    size = lgrp_mem_size(cookie, home, LGRP_MEM_SZ_FREE,
        LGRP_CONTENT_ALL);
    if (size == -1)
        perror("lgrp_mem_size");

    /*
     * It does, so return the home lgroup.
     */
    if (size > 0)
        return (home);

    /*
     * Otherwise, find next nearest lgroup outside of the home.
     */
    return (lgrp_next_nearest(cookie, home));
}

```

EXAMPLE 4-11 Find Nearest lgroup With Free Memory

This example code finds the nearest lgroup with free memory to a given thread's home lgroup.

```

lgrp_id_t
lgrp_nearest(lgrp_cookie_t cookie)
{
    lgrp_id_t    home;
    longlong_t  size;

    /*
     * Get home lgroup
     */

    home = lgrp_home();

    /*
     * See whether home lgroup has any memory available in its hierarchy
     */

    if (lgrp_mem_size(cookie, lgrp, LGRP_MEM_SZ_FREE,
        LGRP_CONTENT_ALL, &size) == -1)
        perror("lgrp_mem_size");

    /*
     * It does, so return the home lgroup.
     */

    if (size > 0)
        return (home);

    /*
     * Otherwise, find next nearest lgroup outside of the home.
     */
}

```

EXAMPLE 4-11 Find Nearest lgroup With Free Memory (Continued)

```
        return (lgrp_next_nearest(cookie, home));  
    }
```

Input/Output Interfaces

This chapter introduces file input/output operations, as provided on systems that do not provide virtual memory services. The chapter discusses the improved input/output method provided by the virtual memory facilities. The chapter describes the older method of locking files and records in [“Using File and Record Locking” on page 84](#).

Files and I/O Interfaces

Files that are organized as a sequence of data are called *regular* files. Regular files can contain ASCII text, text in some other binary data encoding, executable code, or any combination of text, data, and code.

A regular file is made up of the following components:

- Control data, which is called the *inode*. This data includes the file type, the access permissions, the owner, the file size, and the location of the data blocks.
- File contents: a nonterminated sequence of bytes.

The Oracle Solaris operating system provides the following basic forms of file input/output interfaces:

- The traditional, raw style of file I/O is described in [“Basic File I/O” on page 82](#).
- The standard I/O buffering provides an easier interface and improved efficiency to an application run on a system without virtual memory. In an application running in a virtual memory environment, such as on the SunOS operating system, standard file I/O is outdated.
- The memory mapping interface is described in [“Memory Management Interfaces” on page 15](#). Mapping files is the most efficient form of file I/O for most applications run under the SunOS platform.

Basic File I/O

The following interfaces perform basic operations on files and on character I/O devices.

TABLE 5-1 Basic File I/O Interfaces

Interface Name	Purpose
<code>open(2)</code>	Open a file for reading or writing
<code>close(2)</code>	Close a file descriptor
<code>read(2)</code>	Read from a file
<code>write(2)</code>	Write to a file
<code>creat(2)</code>	Create a new file or rewrite an existing one
<code>unlink(2)</code>	Remove a directory entry
<code>lseek(2)</code>	Move read/write file pointer

The following code sample demonstrates the use of the basic file I/O interface. `read(2)` and `write(2)` both transfer no more than the specified number of bytes, starting at the current offset into the file. The number of bytes actually transferred is returned. The end of a file is indicated on a `read(2)` by a return value of zero.

EXAMPLE 5-1 Basic File I/O Interface

```
#include <fcntl.h>
#define MAXSIZE 256

main()
{
    int fd;
    ssize_t n;
    char array[MAXSIZE];

    fd = open ("/etc/motd", O_RDONLY);
    if (fd == -1) {
        perror ("open");
        exit (1);
    }
    while ((n = read (fd, array, MAXSIZE)) > 0)
        if (write (1, array, n) != n)
            perror ("write");
    if (n == -1)
        perror ("read");
    close (fd);
}
```

When you are done reading or writing a file, always call `close(2)`. Do not call `close(2)` for a file descriptor that was not returned from a call to `open(2)`.

File pointer offsets into an open file are changed by using `read(2)`, `write(2)`, or by calls to `lseek(2)`. The following example demonstrates the uses of `lseek`.

```

off_t    start, n;
struct   record  rec;

/* record current offset in start */
start = lseek (fd, 0L, SEEK_CUR);

/* go back to start */
n = lseek (fd, -start, SEEK_SET);
read (fd, &rec, sizeof (rec));

/* rewrite previous record */
n = lseek (fd, -sizeof (rec), SEEK_CUR);
write (fd, (char *)&rec, sizeof (rec));

```

Advanced File I/O

The following table lists the tasks performed by advanced file I/O interfaces.

TABLE 5-2 Advanced File I/O Interfaces

Interface Name	Purpose
<code>link(2)</code> , <code>linkat(2)</code>	Link to a file
<code>access(2)</code> , <code>faccessat(2)</code>	Determine accessibility of a file
<code>mknod(2)</code>	Make a special or ordinary file
<code>chmod(2)</code> , <code>fchmodat(2)</code>	Change mode of file
<code>chown(2)</code> , <code>lchown(2)</code> , <code>fchown(2)</code> , <code>fchownat(2)</code>	Change owner and group of a file
<code>utime(2)</code>	Set file access and modification times
<code>stat(2)</code> , <code>lstat(2)</code> , <code>fstat(2)</code> , <code>fstatat(2)</code>	Get file status
<code>fcntl(2)</code>	Perform file control functions
<code>ioctl(2)</code>	Control device
<code>fpathconf(2)</code>	Get configurable path name variables
<code>opendir(3C)</code> , <code>readdir(3C)</code> , <code>closedir(3C)</code>	Perform directory operations
<code>mkdir(2)</code> , <code>mkdirat(2)</code>	Make a directory
<code>readlink(2)</code> , <code>readlinkat(2)</code>	Read the value of a symbolic link
<code>rename(2)</code> , <code>renameat(2)</code>	Change the name of a file
<code>rmdir(2)</code> , <code>unlinkat(2)</code>	Remove a directory

TABLE 5-2 Advanced File I/O Interfaces (Continued)

Interface Name	Purpose
<code>symlink(2)</code> , <code>symlinkat(2)</code>	Make a symbolic link to a file

See “syscall Provider” in *Oracle Solaris 11.1 Dynamic Tracing Guide* for related information.

File System Control

The file system control interfaces listed in the following table enable the control of various aspects of the file system.

TABLE 5-3 File System Control Interfaces

Interface Name	Purpose
<code>ustat(2)</code>	Get file system statistics
<code>sync(2)</code>	Update super block
<code>mount(2)</code>	Mount a file system
<code>statvfs(2)</code> , <code>fstatvfs(2)</code>	Get file system information
<code>sysfs(2)</code>	Get file system type information

Using File and Record Locking

You do not need to use traditional file I/O to lock file elements. Use the lighter weight synchronization mechanisms that are described in *Multithreaded Programming Guide* with mapped files.

Locking files prevents errors that can occur when several users try to update a file at the same time. You can lock a portion of a file.

File locking blocks access to an entire file. Record locking blocks access to a specified segment of the file. In SunOS, all files are a sequence of bytes of data: a record is a concept of the programs that use the file.

Choosing a Lock Type

Mandatory locking suspends a process until the requested file segments are free. Advisory locking returns a result indicating whether the lock was obtained or not. A process can ignore the result of advisory locking. You cannot use both mandatory and advisory file locking on the same file at the same time. The mode of a file at the time the file is opened determines whether locks on a file are treated as mandatory or advisory.

Of the two basic locking calls, `fcntl(2)` is more portable, more powerful, and less easy to use than `lockf(3C)`. `fcntl(2)` is specified in POSIX 1003.1 standard. `lockf(3C)` is provided to be compatible with older applications.

Selecting Advisory or Mandatory Locking

For mandatory locks, the file must be a regular file with the set-group-ID bit on and the group execute permission off. If either condition fails, all record locks are advisory.

Set a mandatory lock as follows.

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

int mode;
struct stat buf;
...
if (stat(filename, &buf) < 0) {
    perror("program");
    exit (2);
}
/* get currently set mode */
mode = buf.st_mode;
/* remove group execute permission from mode */
mode &= ~(S_IEXEC>>3);
/* set 'set group id bit' in mode */
mode |= S_ISGID;
if (chmod(filename, mode) < 0) {
    perror("program");
    exit(2);
}
...
```

The operating system ignores record locks when the system is executing a file. Any files with record locks should not have execute permissions set.

The `chmod(1)` command can also be used to set a file to permit mandatory locking.

```
$ chmod +l file
```

This command sets the `020n0` permission bit in the file mode, which indicates mandatory locking on the file. If *n* is even, the bit is interpreted as enabling mandatory locking. If *n* is odd, the bit is interpreted as “set group ID on execution.”

The `ls(1)` command shows this setting when you ask for the long listing format with the `-l` option:

```
$ ls -l file
```

This command displays the following information:

```
-rw---l--- 1 user group size mod_time file
```

The letter “l” in the permissions indicates that the set-group-ID bit is on. Since the set-group-ID bit is on, mandatory locking is enabled. Normal semantics of set group ID are also enabled.

Cautions About Mandatory Locking

Keep in mind the following aspects of locking:

- Mandatory locking works only for local files. Mandatory locking is not supported when accessing files through NFS.
- Mandatory locking protects only the segments of a file that are locked. The remainder of the file can be accessed according to normal file permissions.
- If multiple reads or writes are needed for an atomic transaction, the process should explicitly lock all such segments before any I/O begins. Advisory locks are sufficient for all programs that perform in this way.
- Arbitrary programs should not have unrestricted access permission to files on which record locks are used.
- Advisory locking is more efficient because a record lock check does not have to be performed for every I/O request.

Supported File Systems

Both advisory and mandatory locking are supported on the file systems listed in the following table.

TABLE 5-4 Supported File Systems

File System	Description
ufs	The disk-based file system
fifo fs	A pseudo file system of named pipe files that give processes common access to data
name fs	A pseudo file system used mostly by STREAMS for dynamic mounts of file descriptors on top of file
spec fs	A pseudo file system that provides access to special character devices and block devices
zfs	A transactional file system that uses the concept of storage pools to manage physical storage. See <i>Oracle Solaris 11.1 Administration: ZFS File Systems</i> for detailed information.

Only advisory file locking is supported on NFS. File locking is not supported for the `proc` and `fd` file systems.

Opening a File for Locking

You can only request a lock for a file with a valid open descriptor. For read locks, the file must be open with at least read access. For write locks, the file must also be open with write access. In the following example, a file is opened for both read and write access.

```
...
filename = argv[1];
fd = open (filename, O_RDWR);
if (fd < 0) {
    perror(filename);
    exit(2);
}
...
```

Setting a File Lock

To lock an entire file, set the offset to zero and set the size to zero.

You can set a lock on a file in several ways. The choice of method depends on how the lock interacts with the rest of the program, performance, and portability. This example uses the POSIX standard-compatible `fcntl(2)` interface. The interface tries to lock a file until one of the following happens:

- The file lock is set successfully.
- An error occurs.
- `MAX_TRY` is exceeded, and the program stops trying to lock the file.

```
#include <fcntl.h>

...
    struct flock lck;

...
    lck.l_type = F_WRLCK;    /* setting a write lock */
    lck.l_whence = 0;       /* offset l_start from beginning of file */
    lck.l_start = (off_t)0;
    lck.l_len = (off_t)0;   /* until the end of the file */
    if (fcntl(fd, F_SETLK, &lck) < 0) {
        if (errno == EAGAIN || errno == EACCES) {
            (void) fprintf(stderr, "File busy try again later!\n");
            return;
        }
        perror("fcntl");
        exit (2);
    }
    ...
```

Using `fcntl(2)`, you can set the type and start of the lock request by setting structure variables.

Note – You cannot lock mapped files with `flock`. However, you can use the multithread-oriented synchronization mechanisms with mapped files. These synchronization mechanisms can be used in POSIX styles as well as in Oracle Solaris styles.

Setting and Removing Record Locks

When locking a record, do not set the starting point and length of the lock segment to zero. The locking procedure is otherwise identical to file locking.

Contention for data is why you use record locking. Therefore, you should have a failure response for when you cannot obtain all the required locks:

- Wait a certain amount of time, then try again
- Abort the procedure, warn the user
- Let the process sleep until signaled that the lock has been freed
- Do some combination of the previous

This example shows a record being locked by using `fcntl(2)`.

```
{
    struct flock lck;
    ...
    lck.l_type = F_WRLCK;    /* setting a write lock */
    lck.l_whence = 0;      /* offset l_start from beginning of file */
    lck.l_start = here;
    lck.l_len = sizeof(struct record);

    /* lock "this" with write lock */
    lck.l_start = this;
    if (fcntl(fd, F_SETLKW, &lck) < 0) {
        /* "this" lock failed. */
        return (-1);
    }
    ...
}
```

The next example shows the `lockf(3C)` interface.

```
#include <unistd.h>

{
    ...
    /* lock "this" */
    (void) lseek(fd, this, SEEK_SET);
    if (lockf(fd, F_LOCK, sizeof(struct record)) < 0) {
        /* Lock on "this" failed. Clear lock on "here". */
        (void) lseek(fd, here, 0);
        (void) lockf(fd, F_ULOCK, sizeof(struct record));
        return (-1);
    }
}
```

You remove locks in the same way the locks were set. Only the lock type is different (`F_ULOCK`). An unlock cannot be blocked by another process and affects only locks placed by the calling process. The unlock affects only the segment of the file specified in the preceding locking call.

Getting Lock Information

You can determine which process is holding a lock. A lock is set, as in the previous examples, and `F_GETLK` is used in `fcntl(2)`.

The next example finds and prints identifying data on all the locked segments of a file.

EXAMPLE 5-2 Printing Locked Segments of a File

```
struct flock lck;

lck.l_whence = 0;
lck.l_start = 0L;
lck.l_len = 0L;
do {
    lck.l_type = F_WRLCK;
    (void) fcntl(fd, F_GETLK, &lck);
    if (lck.l_type != F_UNLCK) {
        (void) printf("%d %d %c %8ld %8ld\n", lck.l_sysid, lck.l_pid,
            (lck.l_type == F_WRLCK) ? 'W' : 'R', lck.l_start, lck.l_len);
        /* If this lock goes to the end of the address space, no
         * need to look further, so break out. */
        if (lck.l_len == 0) {
            /* else, look for new lock after the one just found. */
            lck.l_start += lck.l_len;
        }
    }
} while (lck.l_type != F_UNLCK);
```

`fcntl(2)` with the `F_GETLK` command can sleep while waiting for a server to respond. The command can fail, returning `ENOLCK`, if either the client or the server have a resource shortage.

Use `lockf(3C)` with the `F_TEST` command to test if a process is holding a lock. This interface does not return information about the lock's location or ownership.

EXAMPLE 5-3 Testing a Process With lockf

```
(void) lseek(fd, 0, 0L);
/* set the size of the test region to zero (0). to test until the
end of the file address space. */
if (lockf(fd, (off_t)0, SEEK_SET) < 0) {
    switch (errno) {
        case EACCES:
        case EAGAIN:
            (void) printf("file is locked by another process\n");
            break;
        case EBADF:
            /* bad argument passed to lockf */
            perror("lockf");
    }
}
```

EXAMPLE 5-3 Testing a Process With `lockf` (Continued)

```

        break;
    default:
        (void) printf("lockf: unexpected error <%d>\n", errno);
        break;
    }
}

```

Process Forking and Locks

When a process forks, the child receives a copy of the file descriptors that the parent opened. Locks are not inherited by the child because the locks are owned by a specific process. The parent and child share a common file pointer for each file. Both processes can try to set locks on the same location in the same file. This problem occurs with both `lockf(3C)` and `fcntl(2)`. If a program holding a record lock forks, the child process should close the file. After closing the file, the child process should reopen the file to set a new, separate file pointer.

Deadlock Handling

The UNIX locking facilities provide deadlock detection and avoidance. Deadlocks can occur only when the system is ready to put a record-locking interface to sleep. A search is made to determine whether two processes are in a deadlock. If a potential deadlock is detected, the locking interface fails and sets `errno` to indicate deadlock. Processes setting locks that use `F_SETLK` do not cause a deadlock because these processes do not wait when the lock cannot be granted immediately.

Terminal I/O Functions

Terminal I/O interfaces deal with a general terminal interface for controlling asynchronous communications ports, as shown in the following table. For more information, see the [termios\(3C\)](#) and [termio\(7I\)](#) man pages.

TABLE 5-5 Terminal I/O Interfaces

Interface Name	Purpose
<code>tcgetattr(3C)</code> , <code>tcsetattr(3C)</code>	Get and set terminal attributes
<code>tcsendbreak(3C)</code> , <code>tcdrain(3C)</code> , <code>tcflush(3C)</code> , <code>tcflow(3C)</code>	Perform line control interfaces
<code>cfgetospeed(3C)</code> , <code>cfgetispeed(3C)</code> , <code>cfsetispeed(3C)</code> , <code>cfsetospeed(3C)</code>	Get and set baud rate
<code>tcsetpgrp(3C)</code>	Get and set terminal foreground process group ID
<code>tcgetsid(3C)</code>	Get terminal session ID

The following example shows how the server dissociates from the controlling terminal of its invoker in the non-DEBUG mode of operation.

EXAMPLE 5-4 Dissociating From the Controlling Terminal

```
(void) close(0);
(void) close(1);
(void) close(2);
(void) open("/", O_RDONLY);
(void) dup2(0, 1);
(void) dup2(0, 2);
setsid();
```

This operation mode prevents the server from receiving signals from the process group of the controlling terminal. A server cannot send reports of errors to a terminal after the server has dissociated. The dissociated server must log errors with `syslog(3C)`.

Interprocess Communication

This chapter is for programmers who develop multiprocess applications.

SunOS 5.11 and compatible operating systems have a large variety of mechanisms for concurrent processes to exchange data and synchronize execution. All of these mechanisms, except mapped memory, are introduced in this chapter.

- Pipes (anonymous data queues) are described in [“Pipes Between Processes” on page 93](#).
- Named pipes (data queues with file names.) [“Named Pipes” on page 95](#) covers named pipes.
- System V message queues, semaphores, and shared memory are described in [“System V IPC” on page 98](#).
- POSIX message queues, semaphores, and shared memory are described in [“POSIX Interprocess Communication” on page 96](#).
- [“Sockets Overview” on page 95](#) describes interprocess communication using sockets.
- Mapped memory and files are described in [“Memory Management Interfaces” on page 15](#).
- Doors (a mechanism for secure control transfer) are described in [“Doors Overview” on page 95](#).

Pipes Between Processes

A pipe between two processes is a pair of files that is created in a parent process. The pipe connects the resulting processes when the parent process forks. A pipe has no existence in any file name space, so it is said to be anonymous. A pipe usually connects only two processes, although any number of child processes can be connected to each other and their related parent by a single pipe.

A pipe is created in the process that becomes the parent by a call to `pipe(2)`. The call returns two file descriptors in the array passed to it. After forking, both processes read from `p[0]` and write to `p[1]`. The processes actually read from and write to a circular buffer that is managed for them.

Because calling `fork(2)` duplicates the per-process open file table, each process has two readers and two writers. Closing the extra readers and writers enables the proper functioning of the pipe. For example, no end-of-file indication would ever be returned if the other end of a reader is left open for writing by the same process. The following code shows pipe creation, a fork, and clearing the duplicate pipe ends.

```
#include <stdio.h>
#include <unistd.h>
...
    int p[2];
...
    if (pipe(p) == -1) exit(1);
    switch( fork() )
    {
        case 0: /* in child */
            close( p[0] );
            dup2( p[1], 1);
            close P[1] );
            exec( ... );
            exit(1);
        default: /* in parent */
            close( p[1] );
            dup2( P[0], 0 );
            close( p[0] );
            break;
    }
    ...
```

The following table shows the results of reads from a pipe and writes to a pipe, under certain conditions.

TABLE 6-1 Read/Write Results in a Pipe

Attempt	Conditions	Result
read	Empty pipe, writer attached	Read blocked
write	Full pipe, reader attached	Write blocked
read	Empty pipe, no writer attached	EOF returned
write	No reader	SIGPIPE

Blocking can be prevented by calling `fcntl(2)` on the descriptor to set `FNDELAY`. This causes an error return (-1) from the I/O call with `errno` set to `EWOULDBLOCK`.

Named Pipes

Named pipes function much like pipes, but are created as named entities in a file system. This enables the pipe to be opened by all processes with no requirement that they be related by forking. A named pipe is created by a call to `mknod(2)`. Any process with appropriate permission can then read or write to a named pipe.

In the `open(2)` call, the process opening the pipe blocks until another process also opens the pipe.

To open a named pipe without blocking, the `open(2)` call joins the `O_NDELAY` mask (found in `sys/fcntl.h`) with the selected file mode mask using the Boolean or operation on the call to `open(2)`. If no other process is connected to the pipe when `open(2)` is called, `-1` is returned with `errno` set to `EWOULDBLOCK`.

Sockets Overview

Sockets provide point-to-point, two-way communication between two processes. Sockets are a basic component of interprocess and intersystem communication. A socket is an endpoint of communication to which a name can be bound. It has a type and one or more associated processes.

Sockets exist in communication domains. A socket domain is an abstraction that provides an addressing structure and a set of protocols. Sockets connect only with sockets in the same domain. Twenty three socket domains are identified (see `sys/socket.h`), of which only the UNIX and Internet domains are normally used in Oracle Solaris OS and compatible operating systems.

You can use sockets to communicate between processes on a single system, like other forms of IPC. The UNIX domain (`AF_UNIX`) provides a socket address space on a single system. UNIX domain sockets are named with UNIX paths. UNIX domain sockets are further described in [Appendix A, “UNIX Domain Sockets.”](#) Sockets can also be used to communicate between processes on different systems. The socket address space between connected systems is called the Internet domain (`AF_INET`). Internet domain communication uses the TCP/IP internet protocol suite. Internet domain sockets are described in [Chapter 7, “Socket Interfaces.”](#)

Doors Overview

Doors are a fast light-weight RPC mechanism for secure control transfer between processes on the same machine. A door is created when a process known as the door server calls `door_create(3DOOR)` with a server function and receives a file descriptor. The file descriptor can be passed to other processes or attached to the file system using `fat_tach(3C)`. A client process, which has the file descriptor, can then invoke the door process by calling

`door_call(3DOOR)`. The client can also pass data and descriptors including other door descriptors. As a result of the call to `door_call()`, the client thread blocks and a thread in the door server wakes up and starts running the server function. When the server function is completed, the function calls `door_return(3DOOR)` to pass optional data and descriptors back to the client. `door_return()` also switches control back to the client; the server thread gets blocked in the kernel and does not return from the `door_return` call.

Doors are described in the doors library [libdoor\(3LIB\)](#).

POSIX Interprocess Communication

POSIX interprocess communication (IPC) is a variation of System V interprocess communication. Like System V objects, POSIX IPC objects have read and write, but not execute, permissions for the owner, the owner's group, and for others. There is no way for the owner of a POSIX IPC object to assign a different owner. POSIX IPC includes the following features:

t

- Messages allow processes to send formatted data streams to arbitrary processes.
- Semaphores allow processes to synchronize execution.
- Shared memory allows processes to share parts of their virtual address space.

Unlike the System V IPC interfaces, the POSIX IPC interfaces are all multithread safe.

POSIX Messages

The POSIX message queue interfaces are listed in the following table.

TABLE 6-2 POSIX Message Queue Interfaces

Interface Name	Purpose
<code>mq_open</code>	Connects to, and optionally creates, a named message queue
<code>mq_close</code>	Ends the connection to an open message queue
<code>mq_unlink</code>	Ends the connection to an open message queue and causes the queue to be removed when the last process closes it
<code>mq_send</code>	Places a message in the queue
<code>mq_receive</code>	Receives (removes) the oldest, highest priority message from the queue
<code>mq_notify</code>	Notifies a process or thread that a message is available in the queue

TABLE 6-2 POSIX Message Queue Interfaces (Continued)

Interface Name	Purpose
mq_setattr	Set or get message queue attributes

POSIX Semaphores

POSIX semaphores are much lighter weight than are System V semaphores. A POSIX semaphore structure defines a single semaphore, not an array of up to 25 semaphores.

The POSIX semaphore interfaces are shown below.

<code>sem_open</code>	Connects to, and optionally creates, a named semaphore
<code>sem_init</code>	Initializes a semaphore structure (internal to the calling program, so not a named semaphore)
<code>sem_close</code>	Ends the connection to an open semaphore
<code>sem_unlink</code>	Ends the connection to an open semaphore and causes the semaphore to be removed when the last process closes it
<code>sem_destroy</code>	Initializes a semaphore structure (internal to the calling program, so not a named semaphore)
<code>sem_getvalue</code>	Copies the value of the semaphore into the specified integer
<code>sem_wait</code>	Blocks while the semaphore is held by other processes or returns an error if the semaphore is held by another process

POSIX Shared Memory

POSIX shared memory is actually a variation of mapped memory (see “[Creating and Using Mappings](#)” on page 15). The major differences are:

- You use `shm_open` to open the shared memory object instead of calling `open(2)`.
- You use `shm_unlink` to close and delete the object instead of calling `close(2)` which does not remove the object.

The options in `shm_open` substantially fewer than the number of options provided in `open(2)`.

System V IPC

SunOS 5.11 and compatible operating systems also provide the System V inter process communication (IPC) package. System V IPC has effectively been replaced by POSIX IPC, but is maintained to support older applications.

See the [ipcrm\(1\)](#), [ipcs\(1\)](#), [Intro\(2\)](#), [msgctl\(2\)](#), [msgget\(2\)](#), [msgrcv\(2\)](#), [msgsnd\(2\)](#), [semget\(2\)](#), [semctl\(2\)](#), [semop\(2\)](#), [shmget\(2\)](#), [shmctl\(2\)](#), [shmop\(2\)](#), and [ftok\(3C\)](#) man pages for more information about System V IPC.

Permissions for Messages, Semaphores, and Shared Memory

Messages, semaphores, and shared memory have read and write permissions, but no execute permission, for the owner, group, and others, which is similar to ordinary files. Like files, the creating process identifies the default owner. Unlike files, the creating process can assign ownership of the facility to another user or revoke an ownership assignment.

IPC Interfaces, Key Arguments, and Creation Flags

Processes requesting access to an IPC facility must be able to identify the facility. To identify the facility to which the process requests access, interfaces that initialize or provide access to an IPC facility use a `key_t` *key* argument. The *key* is an arbitrary value or one that can be derived from a common seed at runtime. One way to derive such a key is by using [ftok\(3C\)](#), which converts a file name to a key value that is unique within the system.

Interfaces that initialize or get access to messages, semaphores, or shared memory return an ID number of type `int`. IPC Interfaces that perform read, write, and control operations use this ID.

If the key argument is specified as `IPC_PRIVATE`, the call initializes a new instance of an IPC facility that is private to the creating process.

When the `IPC_CREAT` flag is supplied in the flags argument appropriate to the call, the interface tries to create the facility if it does not exist already.

When called with both the `IPC_CREAT` and `IPC_EXCL` flags, the interface fails if the facility already exists. This behavior can be useful when more than one process might attempt to initialize the facility. One such case might involve several server processes having access to the same facility. If they all attempt to create the facility with `IPC_EXCL` in effect, only the first attempt succeeds.

If neither of these flags is given and the facility already exists, the interfaces return the ID of the facility to get access. If `IPC_CREAT` is omitted and the facility is not already initialized, the calls fail.

Using logical (bitwise) OR, `IPC_CREAT` and `IPC_EXCL` are combined with the octal permission modes to form the flags argument. For example, the statement below initializes a new message queue if the queue does not exist:

```
msqid = msgget(ftok("/tmp", 'A'), (IPC_CREAT | IPC_EXCL | 0400));
```

The first argument evaluates to a key ('A') based on the string ("/tmp"). The second argument evaluates to the combined permissions and control flags.

System V Messages

Before a process can send or receive a message, you must initialize the queue through `msgget(2)`. The owner or creator of a queue can change its ownership or permissions using `msgctl(2)`. Any process with permission can use `msgctl(2)` for control operations.

IPC messaging enables processes to send and receive messages and queue messages for processing in an arbitrary order. Unlike the file byte-stream data flow of pipes, each IPC message has an explicit length.

Messages can be assigned a specific type. A server process can thus direct message traffic between clients on its queue by using the client process PID as the message type. For single-message transactions, multiple server processes can work in parallel on transactions sent to a shared message queue.

Operations to send and receive messages are performed by `msgsnd(2)` and `msgrcv(2)`, respectively. When a message is sent, its text is copied to the message queue. `msgsnd(2)` and `msgrcv(2)` can be performed as either blocking or non-blocking operations. A blocked message operation remains suspended until one of the following three conditions occurs:

- The call succeeds.
- The process receives a signal.
- The queue is removed.

Initializing a Message Queue

`msgget(2)` initializes a new message queue. It can also return the message queue ID (`msqid`) of the queue corresponding to the key argument. The value passed as the `msgflg` argument must be an octal integer with settings for the queue's permissions and control flags.

The `MSGMNI` kernel configuration option determines the maximum number of unique message queues that the kernel supports. `msgget(2)` fails when this limit is exceeded.

The following code illustrates `msgget(2)`.

```
#include <sys/ipc.h>
#include <sys/msg.h>
...
    key_t    key;          /* key to be passed to msgget() */
    int      msgflg,      /* msgflg to be passed to msgget() */
            msqid;       /* return value from msgget() */
    ...
    key = ...
    msgflg = ...
    if ((msqid = msgget(key, msgflg)) == -1)
    {
        perror("msgget: msgget failed");
        exit(1);
    } else
        (void) fprintf(stderr, "msgget succeeded");
...

```

Controlling Message Queues

`msgctl(2)` alters the permissions and other characteristics of a message queue. The `msqid` argument must be the ID of an existing message queue. The `cmd` argument is one of the following:

- `IPC_STAT` Place information about the status of the queue in the data structure pointed to by `buf`. The process must have read permission for this call to succeed.
- `IPC_SET` Set the owner's user and group ID, the permissions, and the size (in number of bytes) of the message queue. A process must have the effective user ID of the owner, creator, or superuser for this call to succeed.
- `IPC_RMID` Remove the message queue specified by the `msqid` argument.

The following code illustrates `msgctl(2)` with all its various flags.

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>
...
    if (msgctl(msqid, IPC_STAT, &buf) == -1) {
        perror("msgctl: msgctl failed");
        exit(1);
    }
    ...
    if (msgctl(msqid, IPC_SET, &buf) == -1) {
        perror("msgctl: msgctl failed");
        exit(1);
    }
...

```

Sending and Receiving Messages

`msgsnd(2)` and `msgrcv(2)` send and receive messages, respectively. The `msqid` argument must be the ID of an existing message queue. The `msgp` argument is a pointer to a structure that contains the type of the message and its text. The `msgsz` argument specifies the length of the message in bytes. The `msgflg` argument passes various control flags.

The following code illustrates `msgsnd(2)` and `msgrcv(2)`.

```
#include                <sys/types.h>
#include                <sys/ipc.h>
#include                <sys/msg.h>
...
int                    msgflg;        /* message flags for the operation */
struct msgbuf          *msgp;        /* pointer to the message buffer */
size_t                msgsz;         /* message size */
size_t                maxmsgsize;    /* maximum message size */
long                  msgtyp;        /* desired message type */
int                    msqid         /* message queue ID to be used */
...
msgp = malloc(sizeof(struct msgbuf) - sizeof (msgp->mtext)
              + maxmsgsz);
if (msgp == NULL) {
    (void) fprintf(stderr, "msgop: %s %ld byte messages.\n",
                  "could not allocate message buffer for", maxmsgsz);
    exit(1);
    ...
    msgsz = ...
    msgflg = ...
    if (msgsnd(msqid, msgp, msgsz, msgflg) == -1)
        perror("msgop: msgsnd failed");
    ...
    msgsz = ...
    msgtyp = first_on_queue;
    msgflg = ...
    if (rtrn = msgrcv(msqid, msgp, msgsz, msgtyp, msgflg) == -1)
        perror("msgop: msgrcv failed");
    ...
}
```

System V Semaphores

Semaphores enable processes to query or alter status information. They are often used to monitor and control the availability of system resources such as shared memory segments. Semaphores can be operated on as individual units or as elements in a set.

Because System V IPC semaphores can be in a large array, they are extremely heavy weight. Much lighter-weight semaphores are available in the threads library. Also, POSIX semaphores are the most current implementation of System V semaphores (see [“POSIX Semaphores” on page 97](#)). Threads library semaphores must be used with mapped memory (see [“Memory Management Interfaces” on page 15](#)).

A semaphore set consists of a control structure and an array of individual semaphores. A set of semaphores can contain up to 25 elements. The semaphore set must be initialized using `semget(2)`. The semaphore creator can change its ownership or permissions using `semctl(2)`. Any process with permission can use `semctl(2)` to do control operations.

Semaphore operations are performed by `semop(2)`. This interface takes a pointer to an array of semaphore operation structures. Each structure in the array contains data about an operation to perform on a semaphore. Any process with read permission can test whether a semaphore has a zero value. Operations to increment or decrement a semaphore require write permission.

When an operation fails, none of the semaphores are altered. The process blocks unless the `IPC_NOWAIT` flag is set, and remains blocked until:

- The semaphore operations can all finish, so the call succeeds.
- The process receives a signal.
- The semaphore set is removed.

Only one process at a time can update a semaphore. Simultaneous requests by different processes are performed in an arbitrary order. When an array of operations is given by a `semop(2)` call, no updates are done until all operations on the array can finish successfully.

If a process with exclusive use of a semaphore terminates abnormally and fails to undo the operation or free the semaphore, the semaphore stays locked in memory in the state the process left it. To prevent this occurrence, the `SEM_UNDO` control flag makes `semop(2)` allocate an undo structure for each semaphore operation, which contains the operation that returns the semaphore to its previous state. If the process dies, the system applies the operations in the undo structures. This prevents an aborted process from leaving a semaphore set in an inconsistent state.

If processes share access to a resource controlled by a semaphore, operations on the semaphore should not be made with `SEM_UNDO` in effect. If the process that currently has control of the resource terminates abnormally, the resource is presumed to be inconsistent. Another process must be able to recognize this to restore the resource to a consistent state.

When performing a semaphore operation with `SEM_UNDO` in effect, you must also have `SEM_UNDO` in effect for the call that performs the reversing operation. When the process runs normally, the reversing operation updates the undo structure with a complementary value. This ensures that, unless the process is aborted, the values applied to the undo structure are canceled to zero. When the undo structure reaches zero, it is removed.

Using `SEM_UNDO` inconsistently can lead to memory leaks because allocated undo structures might not be freed until the system is rebooted.

Initializing a Semaphore Set

`semget(2)` initializes or gains access to a semaphore. When the call succeeds, it returns the semaphore ID (`semid`). The key argument is a value associated with the semaphore ID. The `nsems` argument specifies the number of elements in a semaphore array. The call fails when `nsems` is greater than the number of elements in an existing array. When the correct count is not known, supplying 0 for this argument ensures that it will succeed. The `semflg` argument specifies the initial access permissions and creation control flags.

The `SEMMNI` system configuration option determines the maximum number of semaphore arrays allowed. The `SEMMNS` option determines the maximum possible number of individual semaphores across all semaphore sets. Because of fragmentation between semaphore sets, allocating all available semaphores might not be possible.

The following code illustrates `semget(2)`.

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>
...
key_t key; /* key to pass to semget() */
int semflg; /* semflg to pass to semget() */
int nsems; /* nsems to pass to semget() */
int semid; /* return value from semget() */
...
key = ...
nsems = ...
semflg = ...
...
if ((semid = semget(key, nsems, semflg)) == -1) {
    perror("semget: semget failed");
    exit(1);
} else
    exit(0);
...
```

Controlling Semaphores

`semctl(2)` changes permissions and other characteristics of a semaphore set. It must be called with a valid semaphore ID. The `semnum` value selects a semaphore within an array by its index. The `cmd` argument is one of the following control flags.

GETVAL	Return the value of a single semaphore.
SETVAL	Set the value of a single semaphore. In this case, <code>arg</code> is taken as <code>arg.val</code> , an <code>int</code> .
GETPID	Return the PID of the process that performed the last operation on the semaphore or array.

GETNCNT	Return the number of processes waiting for the value of a semaphore to increase.
GETZCNT	Return the number of processes waiting for the value of a particular semaphore to reach zero.
GETALL	Return the values for all semaphores in a set. In this case, <code>arg</code> is taken as <code>arg.array</code> , a pointer to an array of unsigned short values.
SETALL	Set values for all semaphores in a set. In this case, <code>arg</code> is taken as <code>arg.array</code> , a pointer to an array of unsigned short values.
IPC_STAT	Return the status information from the control structure for the semaphore set and place it in the data structure pointed to by <code>arg.buf</code> , a pointer to a buffer of type <code>semid_ds</code> .
IPC_SET	Set the effective user and group identification and permissions. In this case, <code>arg</code> is taken as <code>arg.buf</code> .
IPC_RMID	Remove the specified semaphore set.

A process must have an effective user identification of owner, creator, or superuser to perform an `IPC_SET` or `IPC_RMID` command. Read and write permission is required, as for the other control commands.

The following code illustrates `semctl(2)`.

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>
...
    register int i;
...
    i = semctl(semid, semnum, cmd, arg);
    if (i == -1) {
        perror("semctl: semctl failed");
        exit(1);
    }
...
```

Semaphore Operations

`semop(2)` performs operations on a semaphore set. The `semid` argument is the semaphore ID returned by a previous `semget(2)` call. The `sops` argument is a pointer to an array of structures, each containing the following information about a semaphore operation:

- The semaphore number
- The operation to be performed
- Control flags, if any

The `sembuf` structure specifies a semaphore operation, as defined in `sys/sem.h`. The `nsops` argument specifies the length of the array, the maximum size of which is determined by the `SEMOPM` configuration option. This option determines the maximum number of operations allowed by a single `semop(2)` call, and is set to 10 by default.

The operation to be performed is determined as follows:

- Positive integer increments the semaphore value by that amount.
- Negative integer decrements the semaphore value by that amount. An attempt to set a semaphore to a value less than zero fails or blocks, depending on whether `IPC_NOWAIT` is in effect.
- Value of zero means to wait for the semaphore value to reach zero.

The two control flags that can be used with `semop(2)` are `IPC_NOWAIT` and `SEM_UNDO`.

`IPC_NOWAIT` Can be set for any operations in the array. Makes the interface return without changing any semaphore value if it cannot perform any of the operations for which `IPC_NOWAIT` is set. The interface fails if it tries to decrement a semaphore more than its current value, or tests a nonzero semaphore to be equal to zero.

`SEM_UNDO` Allows individual operations in the array to be undone when the process exits.

The following code illustrates `semop(2)`.

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>
...
int i; /* work area */
int nsops; /* number of operations to do */
int semid; /* semid of semaphore set */
struct sembuf *sops; /* ptr to operations to perform */
...
if ((i = semop(semid, sops, nsops)) == -1) {
    perror("semop: semop failed");
} else
    (void) fprintf(stderr, "semop: returned %d\n", i);
...
```

System V Shared Memory

In the SunOS 5.11 operating system, the most efficient way to implement shared memory applications is to rely on `mmap(2)` and on the system's native virtual memory facility. See [Chapter 1, “Memory and CPU Management,”](#) for more information.

The SunOS 5.11 platform also supports System V shared memory, which is a less efficient way to enable the attachment of a segment of physical memory to the virtual address spaces of

multiple processes. When write access is allowed for more than one process, an outside protocol or mechanism, such as a semaphore, can be used to prevent inconsistencies and collisions.

A process creates a shared memory segment using `shmget(2)`. This call is also used to get the ID of an existing shared segment. The creating process sets the permissions and the size in bytes for the segment.

The original owner of a shared memory segment can assign ownership to another user with `shmctl(2)`. The owner can also revoke this assignment. Other processes with proper permission can perform various control functions on the shared memory segment using `shmctl(2)`.

Once created, you can attach a shared segment to a process address space using `shmat(2)`. You can detach it using `shmdt(2)`. The attaching process must have the appropriate permissions for `shmat(2)`. Once attached, the process can read or write to the segment, as allowed by the permission requested in the attach operation. A shared segment can be attached multiple times by the same process.

A shared memory segment is described by a control structure with a unique ID that points to an area of physical memory. The identifier of the segment is called the `shmid`. The structure definition for the shared memory segment control structure can be found in `sys/shm.h`.

Accessing a Shared Memory Segment

`shmget(2)` is used to obtain access to a shared memory segment. When the call succeeds, it returns the shared memory segment ID (`shmid`). The following code illustrates `shmget(2)`.

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>
...
    key_t    key;        /* key to be passed to shmget() */
    int      shmflg;     /* shmflg to be passed to shmget() */
    int      shmid;      /* return value from shmget() */
    size_t   size;      /* size to be passed to shmget() */
    ...
    key = ...
    size = ...
    shmflg = ...
    if ((shmid = shmget (key, size, shmflg)) == -1) {
        perror("shmget: shmget failed");
        exit(1);
    } else {
        (void) fprintf(stderr,
            "shmget: shmget returned %d\n", shmid);
        exit(0);
    }
    ...
```

Controlling a Shared Memory Segment

`shmctl(2)` is used to alter the permissions and other characteristics of a shared memory segment. The `cmd` argument is one of following control commands.

<code>SHM_LOCK</code>	Lock the specified shared memory segment in memory. The process must have the effective ID of superuser to perform this command.
<code>SHM_UNLOCK</code>	Unlock the shared memory segment. The process must have the effective ID of superuser to perform this command.
<code>IPC_STAT</code>	Return the status information contained in the control structure and place it in the buffer pointed to by <code>buf</code> . The process must have read permission on the segment to perform this command.
<code>IPC_SET</code>	Set the effective user and group identification and access permissions. The process must have an effective ID of owner, creator or superuser to perform this command.
<code>IPC_RMID</code>	Remove the shared memory segment. The process must have an effective ID of owner, creator, or superuser to perform this command.

The following code illustrates `shmctl(2)`.

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>
...
int cmd; /* command code for shmctl() */
int shmid; /* segment ID */
struct shm_id_ds shm_id_ds; /* shared memory data structure to hold results */
...
shm_id = ...
cmd = ...
if ((rtrn = shmctl(shmid, cmd, shm_id_ds)) == -1) {
    perror("shmctl: shmctl failed");
    exit(1);
}
...
```

Attaching and Detaching a Shared Memory Segment

`shmat()` and `shmdt()` are used to attach and detach shared memory segments (see the `shmop(2)` man page). `shmat(2)` returns a pointer to the head of the shared segment. `shmdt(2)` detaches the shared memory segment located at the address indicated by `shmaddr`. The following code illustrates calls to `shmat(2)` and `shmdt(2)`

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

```

#include          <sys/shm.h>

static struct state { /* Internal record of attached segments. */
    int          shmidx; /* shmidx of attached segment */
    char         *shmaddr; /* attach point */
    int          shmflg; /* flags used on attach */
} ap[MAXnap];
int nap; /* Number of currently attached segments. */
...
char *addr; /* address work variable */
register int i; /* work area */
register struct state *p; /* ptr to current state entry */
...
    p = &ap[nap++];
    p->shmidx = ...
    p->shmaddr = ...
    p->shmflg = ...
    p->shmaddr = shmat(p->shmidx, p->shmaddr, p->shmflg);
    if(p->shmaddr == (char *)-1) {
        perror("shmat failed");
        nap--;
    } else
        (void) fprintf(stderr, "shmop: shmat returned %p\n",
            p->shmaddr);
    ...
    i = shmdt(addr);
    if(i == -1) {
        perror("shmdt failed");
    } else {
        (void) fprintf(stderr, "shmop: shmdt returned %d\n", i);
        for (p = ap, i = nap; i--; p++) {
            if (p->shmaddr == addr) *p = ap[--nap];
        }
    }
    ...

```

Socket Interfaces

This chapter presents the socket interface. Sample programs are included to illustrate key points. The following topics are discussed in this chapter:

- Socket creation, connection, and closure are discussed in “[Socket Basics](#)” on page 112.
- Client-Server architecture is discussed in “[Client-Server Programs](#)” on page 131.
- Advanced topics such as multicast and asynchronous sockets are discussed in “[Advanced Socket Topics](#)” on page 136.
- Interfaces used to implement the Stream Control Transmission Protocol (SCTP) are discussed in “[Stream Control Transmission Protocol](#)” on page 153.

Note – The interface that is described in this chapter is multithread safe. You can call applications that contain socket interface calls freely in a multithreaded application. Note, however, that the degree of concurrency that is available to applications is not specified.

Overview of Sockets

Sockets have been an integral part of SunOS releases since 1981. A socket is an endpoint of communication to which a name can be bound. A socket has a *type* and an associated process. Sockets were designed to implement the client-server model for interprocess communication where:

- The interface to network protocols needs to accommodate multiple communication protocols, such as TCP/IP, Xerox internet protocols (XNS), and the UNIX family.
- The interface to network protocols needs to accommodate server code that waits for connections and client code that initiates connections.
- Operations differ depending on whether communication is connection-oriented or connectionless.

- Application programs might want to specify the destination address of the datagrams that are being delivered instead of binding the address with the `open(2)` call.

Sockets make network protocols available while behaving like UNIX files. Applications create sockets as sockets are needed. Sockets work with the `close(2)`, `read(2)`, `write(2)`, `ioctl(2)`, and `fcntl(2)` interfaces. The operating system differentiates between the file descriptors for files and the file descriptors for sockets.

Socket Libraries

The socket interface routines are in a library that must be linked with the application. The library `libsocket.so` is contained in `/usr/lib` with the rest of the system service libraries. Use `libsocket.so` for dynamic linking.

Socket Types

Socket types define the communication properties that are visible to a user. The Internet family sockets provide access to the TCP/IP transport protocols. The Internet family is identified by the value `AF_INET6`, for sockets that can communicate over both IPv6 and IPv4. The value `AF_INET` is also supported for source compatibility with old applications and for raw access to IPv4.

The SunOS environment supports four types of sockets:

- *Stream* sockets enable processes to communicate using TCP. A stream socket provides a bidirectional, reliable, sequenced, and unduplicated flow of data with no record boundaries. After the connection has been established, data can be read from and written to these sockets as a byte stream. The socket type is `SOCK_STREAM`.
- *Datagram* sockets enable processes to use UDP to communicate. A datagram socket supports a bidirectional flow of messages. A process on a datagram socket might receive messages in a different order from the sending sequence. A process on a datagram socket might receive duplicate messages. Messages that are sent over a datagram socket might be dropped. Record boundaries in the data are preserved. The socket type is `SOCK_DGRAM`.
- *Raw* sockets provide access to ICMP. Raw sockets also provide access to other protocols based on IP that are not directly supported by the networking stack. These sockets are normally datagram oriented, although their exact characteristics are dependent on the interface provided by the protocol. Raw sockets are not for most applications. Raw sockets are provided to support the development of new communication protocols, or for access to more esoteric facilities of an existing protocol. Only superuser processes can use raw sockets. The socket type is `SOCK_RAW`.
- *SEQ* sockets support 1-to-N Stream Control Transmission Protocol (SCTP) connections. More details on SCTP are in [“Stream Control Transmission Protocol” on page 153](#).

See “Selecting Specific Protocols” on page 140 for further information.

Interface Sets

The SunOS 5.11 platform provides two sets of socket interfaces. The BSD socket interfaces are provided and, since SunOS version 5.7, the XNS 5 (UNIX03) socket interfaces are also provided. The XNS 5 interfaces differ slightly from the BSD interfaces.

The XNS 5 socket interfaces are documented in the following man pages:

- `accept(3XNET)`
- `bind(3XNET)`
- `connect(3XNET)`
- `endhostent(3XNET)`
- `endnetent(3XNET)`
- `endprotoent(3XNET)`
- `endservent(3XNET)`
- `gethostbyaddr(3XNET)`
- `gethostbyname(3XNET)`
- `gethostent(3XNET)`
- `gethostname(3XNET)`
- `getnetbyaddr(3XNET)`
- `getnetbyname(3XNET)`
- `getnetent(3XNET)`
- `getpeername(3XNET)`
- `getprotobyname(3XNET)`
- `getprotobynumber(3XNET)`
- `getprotoent(3XNET)`
- `getservbyname(3XNET)`
- `getservbyport(3XNET)`
- `getservent(3XNET)`
- `getsockname(3XNET)`
- `getsockopt(3XNET)`
- `htonl(3XNET)`
- `htons(3XNET)`
- `inet_addr(3XNET)`
- `inet_lnaof(3XNET)`
- `inet_makeaddr(3XNET)`
- `inet_netof(3XNET)`
- `inet_network(3XNET)`
- `inet_ntoa(3XNET)`
- `listen(3XNET)`
- `ntohl(3XNET)`
- `ntohs(3XNET)`

- `recv(3XNET)`
- `recvfrom(3XNET)`
- `recvmsg(3XNET)`
- `send(3XNET)`
- `sendmsg(3XNET)`
- `sendto(3XNET)`
- `sethostent(3XNET)`
- `setnetent(3XNET)`
- `setprotoent(3XNET)`
- `setservent(3XNET)`
- `setsockopt(3XNET)`
- `shutdown(3XNET)`
- `socket(3XNET)`
- `socketpair(3XNET)`

The traditional BSD Socket behavior is documented in the corresponding 3N man pages. In addition, the following new interfaces have been added to section 3N:

- `freeaddrinfo(3SOCKET)`
- `freehostent(3SOCKET)`
- `getaddrinfo(3SOCKET)`
- `getipnodebyaddr(3SOCKET)`
- `getipnodebyname(3SOCKET)`
- `getnameinfo(3SOCKET)`
- `inet_ntop(3SOCKET)`
- `inet_pton(3SOCKET)`

See the [standards\(5\)](#) man page for information on building applications that use the XNS 5 (UNIX03) socket interface.

Socket Basics

This section describes the use of the basic socket interfaces.

Socket Creation

The `socket(3SOCKET)` call creates a socket in the specified family and of the specified type.

```
s = socket(family, type, protocol);
```

If the protocol is unspecified, the system selects a protocol that supports the requested socket type. The socket handle is returned. The socket handle is a file descriptor.

The *family* is specified by one of the constants that are defined in `sys/socket.h`. Constants that are named `AF_`*suite* specify the address format to use in interpreting names:

AF_APPLETALK	Apple Computer Inc. Appletalk network
AF_INET6	Internet family for IPv6 and IPv4
AF_INET	Internet family for IPv4 only
AF_PUP	Xerox Corporation PUP internet
AF_UNIX	UNIX file system

Socket types are defined in `sys/socket.h`. These types, `SOCK_STREAM`, `SOCK_DGRAM`, or `SOCK_RAW`, are supported by `AF_INET6`, `AF_INET`, and `AF_UNIX`. The following example creates a stream socket in the Internet family:

```
s = socket(AF_INET6, SOCK_STREAM, 0);
```

This call results in a stream socket. The TCP protocol provides the underlying communication. Set the *protocol* argument to `0`, the default, in most situations. You can specify a protocol other than the default, as described in [“Advanced Socket Topics” on page 136](#).

Binding Local Names

A socket is created with no name. A remote process has no way to refer to a socket until an address is bound to the socket. Processes that communicate are connected through addresses. In the Internet family, a connection is composed of local and remote addresses and local and remote ports. Duplicate ordered sets, such as: protocol, local address, local port, foreign address, foreign port cannot exist. In most families, connections must be unique.

The `bind(3SOCKET)` interface enables a process to specify the local address of the socket. This interface forms the local address, local port set. `connect(3SOCKET)` and `accept(3SOCKET)` complete a socket's association by fixing the remote half of the address tuple. The `bind(3SOCKET)` call is used as follows:

```
bind (s, name, namelen);
```

The socket handle is `s`. The bound name is a byte string that is interpreted by the supporting protocols. Internet family names contain an Internet address and port number.

This example demonstrates binding an Internet address.

```
#include <sys/types.h>
#include <netinet/in.h>
...
struct sockaddr_in6 sin6;
...
s = socket(AF_INET6, SOCK_STREAM, 0);
bzero (&sin6, sizeof (sin6));
sin6.sin6_family = AF_INET6;
```

```
sin6.sin6_addr.s6_addr = in6addr_arg;
sin6.sin6_port = htons(MYPORT);
bind(s, (struct sockaddr *) &sin6, sizeof sin6);
```

The content of the address `sin6` is described in [“Address Binding” on page 140](#), where Internet address bindings are discussed.

Connection Establishment

Connection establishment is usually asymmetric, with one process acting as the client and the other as the server. The server binds a socket to a well-known address associated with the service and blocks on its socket for a connect request. An unrelated process can then connect to the server. The client requests services from the server by initiating a connection to the server's socket. On the client side, the `connect(3SOCKET)` call initiates a connection. In the Internet family, this connection might appear as:

```
struct sockaddr_in6 server;
...
connect(s, (struct sockaddr *)&server, sizeof server);
```

If the client's socket is unbound at the time of the connect call, the system automatically selects and binds a name to the socket. For more information, see [“Address Binding” on page 140](#). This automatic selection is the usual way to bind local addresses to a socket on the client side.

To receive a client's connection, a server must perform two steps after binding its socket. The first step is to indicate how many connection requests can be queued. The second step is to accept a connection.

```
struct sockaddr_in6 from;
...
listen(s, 5);           /* Allow queue of 5 connections */
fromlen = sizeof(from);
newsock = accept(s, (struct sockaddr *) &from, &fromlen);
```

The socket handle `s` is the socket bound to the address to which the connection request is sent. The second parameter of `listen(3SOCKET)` specifies the maximum number of outstanding connections that might be queued. The `from` structure is filled with the address of the client. A NULL pointer might be passed. `fromlen` is the length of the structure.

The `accept(3SOCKET)` routine normally blocks processes. `accept(3SOCKET)` returns a new socket descriptor that is connected to the requesting client. The value of `fromlen` is changed to the actual size of the address.

A server cannot indicate that the server accepts connections from only specific addresses. The server can check the `from` address returned by `accept(3SOCKET)` and close a connection with

an unacceptable client. A server can accept connections on more than one socket, or avoid blocking on the `accept(3SOCKET)` call. These techniques are presented in “Advanced Socket Topics” on page 136.

Connection Errors

An error is returned if the connection is unsuccessful, but an address bound by the system remains. If the connection is successful, the socket is associated with the server and data transfer can begin.

The following table lists some of the more common errors returned when a connection attempt fails.

TABLE 7-1 Socket Connection Errors

Socket Errors	Error Description
ENOBUFS	Lack of memory available to support the call.
EPROTONOSUPPORT	Request for an unknown protocol.
EPROTOTYPE	Request for an unsupported type of socket.
ETIMEDOUT	No connection established in specified time. This error happens when the destination host is down or when problems in the network cause in lost transmissions.
ECONNREFUSED	The host refused service. This error happens when a server process is not present at the requested address.
ENETDOWN or EHOSTDOWN	These errors are caused by status information delivered by the underlying communication interface.
ENETUNREACH or EHOSTUNREACH	These operational errors can occur because no route to the network or host exists. These errors can also occur because of status information returned by intermediate gateways or switching nodes. The status information that is returned is not always sufficient to distinguish between a network that is down and a host that is down.

Data Transfer

This section describes the interfaces to send and receive data. You can send or receive a message with the normal `read(2)` and `write(2)` interfaces:

```
write(s, buf, sizeof buf);
read(s, buf, sizeof buf);
```

You can also use `send(3SOCKET)` and `recv(3SOCKET)`:

```
send(s, buf, sizeof buf, flags);  
recv(s, buf, sizeof buf, flags);
```

`send(3SOCKET)` and `recv(3SOCKET)` are very similar to `read(2)` and `write(2)`, but the `flags` argument is important. The `flags` argument, which is defined in `sys/socket.h`, can be specified as a nonzero value if one or more of the following is required:

<code>MSG_OOB</code>	Send and receive out-of-band data
<code>MSG_PEEK</code>	Look at data without reading
<code>MSG_DONTROUTE</code>	Send data without routing packets

Out-of-band data is specific to stream sockets. When `MSG_PEEK` is specified with a `recv(3SOCKET)` call, any data present is returned to the user, but treated as still unread. The next `read(2)` or `recv(3SOCKET)` call on the socket returns the same data. The option to send data without routing packets applied to the outgoing packets is currently used only by the routing table management process.

Closing Sockets

A `SOCK_STREAM` socket can be discarded by a `close(2)` interface call. If data is queued to a socket that promises reliable delivery after a `close(2)`, the protocol continues to try to transfer the data. The data is discarded if it remains undelivered after an arbitrary period.

A `shutdown(3SOCKET)` closes `SOCK_STREAM` sockets gracefully. Both processes can acknowledge that they are no longer sending. This call has the form:

```
shutdown(s, how);
```

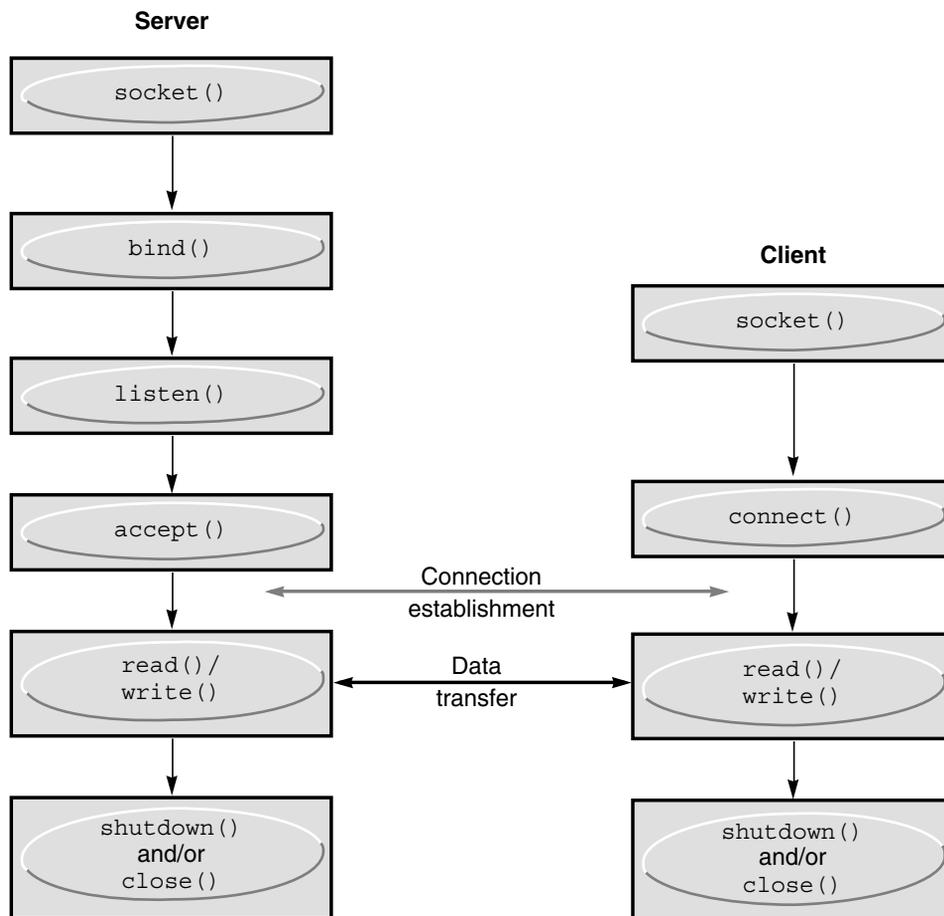
where `how` is defined as

0	Disallows further data reception
1	Disallows further data transmission
2	Disallows further transmission and further reception

Connecting Stream Sockets

The following two examples illustrate initiating and accepting an Internet family stream connection.

FIGURE 7-1 Connection-Oriented Communication Using Stream Sockets



The following example program is a server. The server creates a socket and binds a name to the socket, then displays the port number. The program calls `listen(3SOCKET)` to mark the socket as ready to accept connection requests and to initialize a queue for the requests. The rest of the program is an infinite loop. Each pass of the loop accepts a new connection and removes it from the queue, creating a new socket. The server reads and displays the messages from the socket and closes the socket. The use of `in6addr_any` is explained in “Address Binding” on page 140.

EXAMPLE 7-1 Accepting an Internet Stream Connection (Server)

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

EXAMPLE 7-1 Accepting an Internet Stream Connection (Server) (Continued)

```
#include <stdio.h>
#define TRUE 1
/*
 * This program creates a socket and then begins an infinite loop.
 * Each time through the loop it accepts a connection and prints
 * data from it. When the connection breaks, or the client closes
 * the connection, the program accepts a new connection.
 */
main() {
    int sock, length;
    struct sockaddr_in6 server;
    int msgsock;
    char buf[1024];
    int rval;
    /* Create socket. */
    sock = socket(AF_INET6, SOCK_STREAM, 0);
    if (sock == -1) {
        perror("opening stream socket");
        exit(1);
    }
    /* Bind socket using wildcards.*/
    bzero (&server, sizeof(server));
    server.sin6_family = AF_INET6;
    server.sin6_addr = in6addr_any;
    server.sin6_port = 0;
    if (bind(sock, (struct sockaddr *) &server, sizeof server)
        == -1) {
        perror("binding stream socket");
        exit(1);
    }
    /* Find out assigned port number and print it out. */
    length = sizeof server;
    if (getsockname(sock, (struct sockaddr *) &server, &length)
        == -1) {
        perror("getting socket name");
        exit(1);
    }
    printf("Socket port #d\n", ntohs(server.sin6_port));
    /* Start accepting connections. */
    listen(sock, 5);
    do {
        msgsock = accept(sock, (struct sockaddr *) 0, (int *) 0);
        if (msgsock == -1)
            perror("accept");
        else do {
            memset(buf, 0, sizeof buf);
            if ((rval = read(msgsock, buf, sizeof(buf))) == -1)
                perror("reading stream message");
            if (rval == 0)
                printf("Ending connection\n");
            else
                /* assumes the data is printable */
                printf("-->%s\n", buf);
        } while (rval > 0);
        close(msgsock);
    }
```

EXAMPLE 7-1 Accepting an Internet Stream Connection (Server) *(Continued)*

```

    } while(TRUE);
    /*
     * Since this program has an infinite loop, the socket "sock" is
     * never explicitly closed. However, all sockets are closed
     * automatically when a process is killed or terminates normally.
     */
    exit(0);
}

```

To initiate a connection, the client program in [Example 7-2](#) creates a stream socket, then calls `connect(3SOCKET)`, specifying the address of the socket for connection. If the target socket exists, and the request is accepted, the connection is complete. The program can now send data. Data is delivered in sequence with no message boundaries. The connection is destroyed when either socket is closed. For more information about data representation routines in this program, such as `ntohl(3SOCKET)`, `ntohs(3SOCKET)`, `htons(3SOCKET)`, and `htonl(3XNET)`, see the `byteorder(3SOCKET)` man page.

EXAMPLE 7-2 Internet Family Stream Connection (Client)

```

#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netdb.h>
#include <stdio.h>
#define DATA "Half a league, half a league . . ."
/*
 * This program creates a socket and initiates a connection with
 * the socket given in the command line. Some data are sent over the
 * connection and then the socket is closed, ending the connection.
 * The form of the command line is: streamwrite hostname portnumber
 * Usage: pgm host port
 */
main(int argc, char *argv[])
{
    int sock, errnum, h_addr_index;
    struct sockaddr_in6 server;
    struct hostent *hp;
    char buf[1024];
    /* Create socket. */
    sock = socket( AF_INET6, SOCK_STREAM, 0);
    if (sock == -1) {
        perror("opening stream socket");
        exit(1);
    }
    /* Connect socket using name specified by command line. */
    bzero (&server, sizeof (server));
    server.sin6_family = AF_INET6;
    hp = getipnodebyname(argv[1], AF_INET6, AI_DEFAULT, &errnum);
    /*
     * getipnodebyname returns a structure including the network address
     * of the specified host.
     */
}

```

EXAMPLE 7-2 Internet Family Stream Connection (Client) *(Continued)*

```

if (hp == (struct hostent *) 0) {
    fprintf(stderr, "%s: unknown host\n", argv[1]);
    exit(2);
}
h_addr_index = 0;
while (hp->h_addr_list[h_addr_index] != NULL) {
    bcopy(hp->h_addr_list[h_addr_index], &server.sin6_addr,
          hp->h_length);
    server.sin6_port = htons(atoi(argv[2]));
    if (connect(sock, (struct sockaddr *) &server,
               sizeof (server)) == -1) {
        if (hp->h_addr_list[++h_addr_index] != NULL) {
            /* Try next address */
            continue;
        }
        perror("connecting stream socket");
        freehostent(hp);
        exit(1);
    }
    break;
}
freehostent(hp);
if (write(sock, DATA, sizeof DATA) == -1)
    perror("writing on stream socket");
close(sock);
freehostent(hp);
exit(0);
}

```

You can add support for one-to-one SCTP connections to stream sockets. The following example code adds the `-p` to an existing program, enabling the program to specify the protocol to use.

EXAMPLE 7-3 Adding SCTP Support to a Stream Socket

```

#include <stdio.h>
#include <netdb.h>
#include <string.h>
#include <errno.h>

int
main(int argc, char *argv[])
{
    struct protoent *proto = NULL;
    int c;
    int s;
    int protocol;

    while ((c = getopt(argc, argv, "p:")) != -1) {
        switch (c) {
            case 'p':
                proto = getprotobyname(optarg);
                if (proto == NULL) {

```

EXAMPLE 7-3 Adding SCTP Support to a Stream Socket (Continued)

```

        fprintf(stderr, "Unknown protocol: %s\n",
                optarg);
        return (-1);
    }
    break;
default:
    fprintf(stderr, "Unknown option: %c\n", c);
    return (-1);
}
}

/* Use the default protocol, which is TCP, if not specified. */
if (proto == NULL)
    protocol = 0;
else
    protocol = proto->p_proto;

/* Create a IPv6 SOCK_STREAM socket of the protocol. */
if ((s = socket(AF_INET6, SOCK_STREAM, protocol)) == -1) {
    fprintf(stderr, "Cannot create SOCK_STREAM socket of type %s: "
            "%s\n", proto != NULL ? proto->p_name : "tcp",
            strerror(errno));
    return (-1);
}
printf("Success\n");
return (0);
}

```

Input/Output Multiplexing

Requests can be multiplexed among multiple sockets or multiple files. Use [select\(3C\)](#) to multiplex:

```

#include <sys/time.h>
#include <sys/types.h>
#include <sys/select.h>
...
fd_set readmask, writemask, exceptmask;
struct timeval timeout;
...
select(nfds, &readmask, &writemask, &exceptmask, &timeout);

```

The first argument of [select\(3C\)](#) is the number of file descriptors in the lists pointed to by the next three arguments.

The second, third, and fourth arguments of [select\(3C\)](#) point to three sets of file descriptors: a set of descriptors to read on, a set to write on, and a set on which exception conditions are accepted. Out-of-band data is the only exceptional condition. You can designate any of these pointers as a properly cast null. Each set is a structure that contains an array of long integer bit

masks. Set the size of the array with `FD_SETSIZE`, which is defined in `select.h`. The array is long enough to hold one bit for each `FD_SETSIZE` file descriptor.

The macros `FD_SET(fd, &mask)` and `FD_CLR(fd, &mask)` add and delete, respectively, the file descriptor *fd* in the set *mask*. The set should be zeroed before use and the macro `FD_ZERO(&mask)` clears the set *mask*.

The fifth argument of `select(3C)` enables the specification of a timeout value. If the `timeout` pointer is `NULL`, `select(3C)` blocks until a descriptor is selectable, or until a signal is received. If the fields in `timeout` are set to `0`, `select(3C)` polls and returns immediately.

The `select(3C)` routine normally returns the number of file descriptors that are selected, or a zero if the timeout has expired. The `select(3C)` routine returns `-1` for an error or interrupt, with the error number in `errno` and the file descriptor masks unchanged. For a successful return, the three sets indicate which file descriptors are ready to be read from, written to, or have exceptional conditions pending.

Test the status of a file descriptor in a select mask with the `FD_ISSET(fd, &mask)` macro. The macro returns a nonzero value if *fd* is in the set *mask*. Otherwise, the macro returns zero. Use `select(3C)` followed by a `FD_ISSET(fd, &mask)` macro on the read set to check for queued connect requests on a socket.

The following example shows how to select on a listening socket for readability to determine when a new connection can be picked up with a call to `accept(3SOCKET)`. The program accepts connection requests, reads data, and disconnects on a single socket.

EXAMPLE 7-4 Using `select(3C)` to Check for Pending Connections

```
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/time.h>
#include <netinet/in.h>
#include <netdb.h>
#include <stdio.h>
#define TRUE 1
/*
 * This program uses select to check that someone is
 * trying to connect before calling accept.
 */
main() {
    int sock, length;
    struct sockaddr_in6 server;
    int msgsock;
    char buf[1024];
    int rval;
    fd_set ready;
    struct timeval to;
    /* Open a socket and bind it as in previous examples. */
    /* Start accepting connections. */
    listen(sock, 5);
    do {
```

EXAMPLE 7-4 Using `select(3C)` to Check for Pending Connections (Continued)

```

    FD_ZERO(&ready);
    FD_SET(sock, &ready);
    to.tv_sec = 5;
    to.tv_usec = 0;
    if (select(sock + 1, &ready, (fd_set *)0,
              (fd_set *)0, &to) == -1) {
        perror("select");
        continue;
    }
    if (FD_ISSET(sock, &ready)) {
        msgsock = accept(sock, (struct sockaddr *)0, (int *)0);
        if (msgsock == -1)
            perror("accept");
        else do {
            memset(buf, 0, sizeof buf);
            if ((rval = read(msgsock, buf, sizeof(buf))) == -1)
                perror("reading stream message");
            else if (rval == 0)
                printf("Ending connection\n");
            else
                printf("-->%s\n", buf);
        } while (rval > 0);
        close(msgsock);
    } else
        printf("Do something else\n");
    } while (TRUE);
    exit(0);
}

```

In previous versions of the `select(3C)` routine, its arguments were pointers to integers instead of pointers to `fd_sets`. This style of call still works if the number of file descriptors is smaller than the number of bits in an integer.

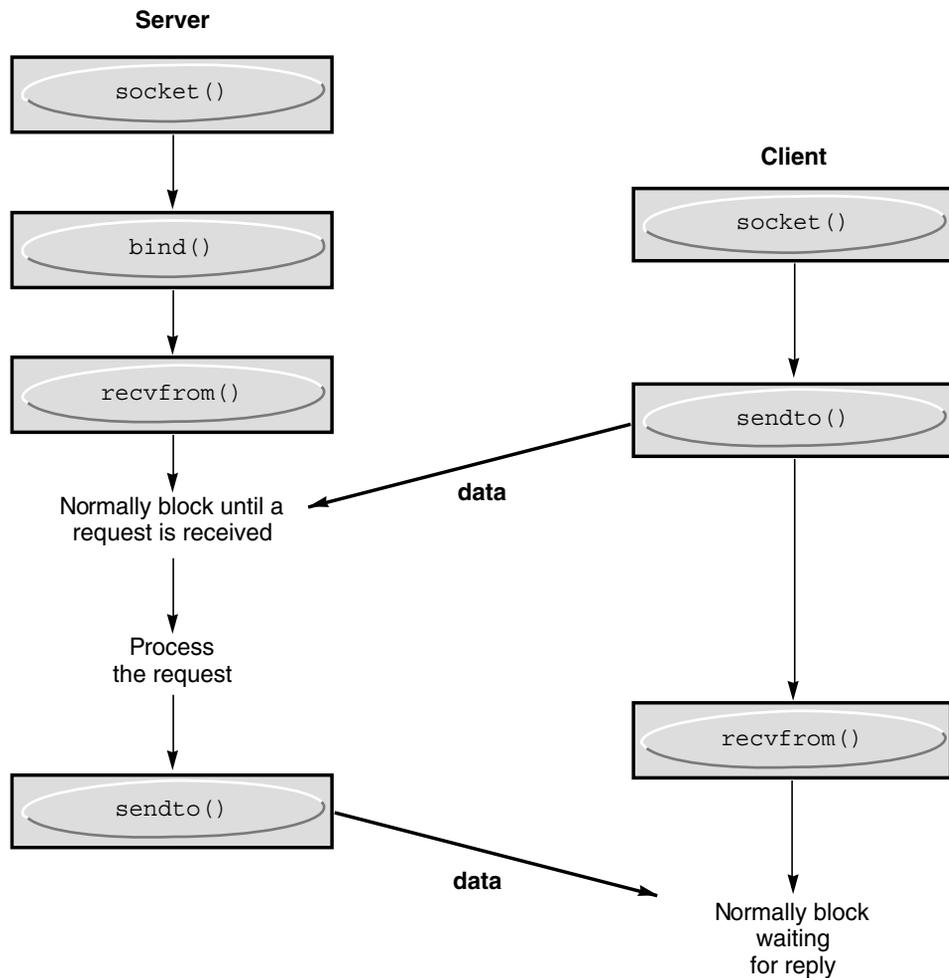
The `select(3C)` routine provides a synchronous multiplexing scheme. The `SIGIO` and `SIGURG` signals, which is described in “[Advanced Socket Topics](#)” on page 136, provide asynchronous notification of output completion, input availability, and exceptional conditions.

Datagram Sockets

A datagram socket provides a symmetric data exchange interface without requiring connection establishment. Each message carries the destination address. The following figure shows the flow of communication between server and client.

The `bind(3SOCKET)` step for the server is optional.

FIGURE 7-2 Connectionless Communication Using Datagram Sockets



Create datagram sockets as described in “[Socket Creation](#)” on page 112. If a particular local address is needed, the `bind(3SOCKET)` operation must precede the first data transmission. Otherwise, the system sets the local address or port when data is first sent. Use `sendto(3SOCKET)` to send data.

```
sendto(s, buf, buflen, flags, (struct sockaddr *) &to, tolen);
```

The *s*, *buf*, *buflen*, and *flags* parameters are the same as in connection-oriented sockets. The *to* and *tolen* values indicate the address of the intended recipient of the message. A locally detected error condition, such as an unreachable network, causes a return of `-1` and *errno* to be set to the error number.

```
recvfrom(s, buf, buflen, flags, (struct sockaddr *) &from, &fromlen);
```

To receive messages on a datagram socket, `recvfrom(3SOCKET)` is used. Before the call, *fromlen* is set to the size of the *from* buffer. On return, *fromlen* is set to the size of the address from which the datagram was received.

Datagram sockets can also use the `connect(3SOCKET)` call to associate a socket with a specific destination address. The socket can then use the `send(3SOCKET)` call. Any data that is sent on the socket that does not explicitly specify a destination address is addressed to the connected peer. Only the data that is received from that peer is delivered. A socket can have only one connected address at a time. A second `connect(3SOCKET)` call changes the destination address. Connect requests on datagram sockets return immediately. The system records the peer's address. Neither `accept(3SOCKET)` nor `listen(3SOCKET)` are used with datagram sockets.

A datagram socket can return errors from previous `send(3SOCKET)` calls asynchronously while the socket is connected. The socket can report these errors on subsequent socket operations. Alternately, the socket can use an option of `getsockopt(3SOCKET)`, `SO_ERROR` to interrogate the error status.

The following example code shows how to send an Internet call by creating a socket, binding a name to the socket, and sending the message to the socket.

EXAMPLE 7-5 Sending an Internet Family Datagram

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netdb.h>
#include <stdio.h>
#define DATA "The sea is calm, the tide is full . . ."
/*
 * Here I send a datagram to a receiver whose name I get from
 * the command line arguments. The form of the command line is:
 * dgramsend hostname portnumber
 */
main(int argc, char *argv[])
{
    int sock, errnum;
    struct sockaddr_in6 name;
    struct hostent *hp;
    /* Create socket on which to send. */
    sock = socket(AF_INET6, SOCK_DGRAM, 0);
    if (sock == -1) {
        perror("opening datagram socket");
        exit(1);
    }
    /*
     * Construct name, with no wildcards, of the socket to "send"
     * to. getinodbyname returns a structure including the network
     * address of the specified host. The port number is taken from
     * the command line.
     */
```

EXAMPLE 7-5 Sending an Internet Family Datagram (Continued)

```
    */
    hp = getipnodebyname(argv[1], AF_INET6, AI_DEFAULT, &errnum);
    if (hp == (struct hostent *) 0) {
        fprintf(stderr, "%s: unknown host\n", argv[1]);
        exit(2);
    }
    bzero (&name, sizeof (name));
    memcpy((char *) &name.sin6_addr, (char *) hp->h_addr,
           hp->h_length);
    name.sin6_family = AF_INET6;
    name.sin6_port = htons(atoi(argv[2]));
    /* Send message. */
    if (sendto(sock, DATA, sizeof DATA, 0,
              (struct sockaddr *) &name, sizeof name) == -1)
        perror("sending datagram message");
    close(sock);
    exit(0);
}
```

The following sample code shows how to read an Internet call by creating a socket, binding a name to the socket, and then reading from the socket.

EXAMPLE 7-6 Reading Internet Family Datagrams

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <stdio.h>
/*
 * This program creates a datagram socket, binds a name to it, then
 * reads from the socket.
 */
main()
{
    int sock, length;
    struct sockaddr_in6 name;
    char buf[1024];
    /* Create socket from which to read. */
    sock = socket(AF_INET6, SOCK_DGRAM, 0);
    if (sock == -1) {
        perror("opening datagram socket");
        exit(1);
    }
    /* Create name with wildcards. */
    bzero (&name, sizeof (name));
    name.sin6_family = AF_INET6;
    name.sin6_addr = in6addr_any;
    name.sin6_port = 0;
    if (bind (sock, (struct sockaddr *)&name, sizeof (name)) == -1) {
        perror("binding datagram socket");
        exit(1);
    }
    /* Find assigned port value and print it out. */
```

EXAMPLE 7-6 Reading Internet Family Datagrams (Continued)

```

length = sizeof(name);
if (getsockname(sock, (struct sockaddr *) &name, &length)
    == -1) {
    perror("getting socket name");
    exit(1);
}
printf("Socket port %#d\n", ntohs(name.sin6_port));
/* Read from the socket. */
if (read(sock, buf, 1024) == -1)
    perror("receiving datagram packet");
/* Assumes the data is printable */
printf("->%s\n", buf);
close(sock);
exit(0);
}

```

Standard Routines

This section describes the routines that you can use to locate and construct network addresses. Unless otherwise stated, interfaces presented in this section apply only to the Internet family.

Locating a service on a remote host requires many levels of mapping before the client and server communicate. A service has a name for human use. The service and host names must translate to network addresses. Finally, the network address must be usable to locate and route to the host. The specifics of the mappings can vary between network architectures.

Standard routines map host names to network addresses, network names to network numbers, protocol names to protocol numbers, and service names to port numbers. The standard routines also indicate the appropriate protocol to use in communicating with the server process. The file `netdb.h` must be included when using any of these routines.

Host and Service Names

The interfaces `getaddrinfo(3SOCKET)`, `getnameinfo(3SOCKET)`, `gai_strerror(3SOCKET)`, and `freeaddrinfo(3SOCKET)` provide a simplified way to translate between the names and addresses of a service on a host. These interfaces are more recent than the `getipnodebyname(3SOCKET)`, `gethostbyname(3NSL)`, and `getservbyname(3SOCKET)` APIs. Both IPv6 and IPv4 addresses are handled transparently.

The `getaddrinfo(3SOCKET)` routine returns the combined address and port number of the specified host and service names. Because the information returned by `getaddrinfo(3SOCKET)` is dynamically allocated, the information must be freed by `freeaddrinfo(3SOCKET)` to prevent memory leaks. `getnameinfo(3SOCKET)` returns the host

and services names associated with a specified address and port number. Call [gai_strerror\(3SOCKET\)](#) to print error messages based on the `EAI_XXX` codes returned by [getaddrinfo\(3SOCKET\)](#) and [getnameinfo\(3SOCKET\)](#).

An example of using [getaddrinfo\(3SOCKET\)](#) follows.

```
struct addrinfo      *res, *aip;
struct addrinfo      hints;
int                  error;

/* Get host address. Any type of address will do. */
bzero(&hints, sizeof (hints));
hints.ai_flags = AI_ALL|AI_ADDRCONFIG;
hints.ai_socktype = SOCK_STREAM;

error = getaddrinfo(hostname, servicename, &hints, &res);
if (error != 0) {
    (void) fprintf(stderr, "getaddrinfo: %s for host %s service %s\n",
        gai_strerror(error), hostname, servicename);
    return (-1);
}
```

After processing the information returned by [getaddrinfo\(3SOCKET\)](#) in the structure pointed to by `res`, the storage should be released by `freeaddrinfo(res)`.

The [getnameinfo\(3SOCKET\)](#) routine is particularly useful in identifying the cause of an error, as in the following example:

```
struct sockaddr_storage faddr;
int sock, new_sock, sock_opt;
socklen_t faddrrlen;
int error;
char hname[NI_MAXHOST];
char sname[NI_MAXSERV];
...
faddrrlen = sizeof (faddr);
new_sock = accept(sock, (struct sockaddr *)&faddr, &faddrrlen);
if (new_sock == -1) {
    if (errno != EINTR && errno != ECONNABORTED) {
        perror("accept");
    }
    continue;
}
error = getnameinfo((struct sockaddr *)&faddr, faddrrlen, hname,
    sizeof (hname), sname, sizeof (sname), 0);
if (error) {
    (void) fprintf(stderr, "getnameinfo: %s\n",
        gai_strerror(error));
} else {
    (void) printf("Connection from %s/%s\n", hname, sname);
}
```

Host Names – hostent

An Internet host-name-to-address mapping is represented by the `hostent` structure as defined in [gethostent\(3NSL\)](#):

```
struct hostent {
    char *h_name;           /* official name of host */
    char **h_aliases;      /* alias list */
    int h_addrtype;        /* hostaddrtype(e.g., AF_INET6) */
    int h_length;          /* length of address */
    char **h_addr_list;    /* list of addrs, null terminated */
};
/*1st addr, net byte order*/
#define h_addr h_addr_list[0]
```

getipnodebyname(3SOCKET)	Maps an Internet host name to a <code>hostent</code> structure
getipnodebyaddr(3SOCKET)	Maps an Internet host address to a <code>hostent</code> structure
freehostent(3SOCKET)	Frees the memory of a <code>hostent</code> structure
inet_ntop(3SOCKET)	Maps an Internet host address to a string

The routines return a `hostent` structure that contains the name of the host, its aliases, the address type, and a NULL-terminated list of variable length addresses. The list of addresses is required because a host can have many addresses. The `h_addr` definition is for backward compatibility, and is the first address in the list of addresses in the `hostent` structure.

Network Names – netent

The routines to map network names to numbers and the reverse return a `netent` structure:

```
/*
 * Assumes that a network number fits in 32 bits.
 */
struct netent {
    char *n_name;          /* official name of net */
    char **n_aliases;     /* alias list */
    int n_addrtype;        /* net address type */
    int n_net;             /* net number, host byte order */
};
```

[getnetbyname\(3SOCKET\)](#), [getnetbyaddr_r\(3SOCKET\)](#), and [getnetent\(3SOCKET\)](#) are the network counterparts to the host routines previously described.

Protocol Names – protoent

The protoent structure defines the protocol-name mapping used with [getprotobyname\(3SOCKET\)](#), [getprotobynumber\(3SOCKET\)](#), and [getprotoent\(3SOCKET\)](#) and defined in [getprotoent\(3SOCKET\)](#):

```
struct protoent {
    char    *p_name;        /* official protocol name */
    char    **p_aliases;    /* alias list */
    int     p_proto;        /* protocol number */
};
```

Service Names – servent

An Internet family service resides at a specific, well-known port, and uses a particular protocol. A service-name-to-port-number mapping is described by the servent structure that is defined in [getprotoent\(3SOCKET\)](#):

```
struct servent {
    char    *s_name;        /* official service name */
    char    **s_aliases;    /* alias list */
    int     s_port;        /* port number, network byte order */
    char    *s_proto;        /* protocol to use */
};
```

[getservbyname\(3SOCKET\)](#) maps service names and, optionally, a qualifying protocol to a servent structure. The call:

```
sp = getservbyname("telnet", (char *) 0);
```

returns the service specification of a telnet server that is using any protocol. The call:

```
sp = getservbyname("telnet", "tcp");
```

returns the telnet server that uses the TCP protocol. [getservbyport\(3SOCKET\)](#) and [getservent\(3SOCKET\)](#) are also provided. [getservbyport\(3SOCKET\)](#) has an interface that is similar to the interface used by [getservbyname\(3SOCKET\)](#). You can specify an optional protocol name to qualify lookups.

Other Routines

Several other routines that simplify manipulating names and addresses are available. The following table summarizes the routines for manipulating variable-length byte strings and byte-swapping network addresses and values.

TABLE 7-2 Runtime Library Routines

Interface	Synopsis
<code>memcmp(3C)</code>	Compares byte-strings; 0 if same, not 0 otherwise
<code>memcpy(3C)</code>	Copies <i>n</i> bytes from <i>s2</i> to <i>s1</i>
<code>memset(3C)</code>	Sets <i>n</i> bytes to <i>value</i> starting at <i>base</i>
<code>htonl(3SOCKET)</code>	32-bit quantity from host into network byte order
<code>htons(3SOCKET)</code>	16-bit quantity from host into network byte order
<code>ntohl(3SOCKET)</code>	32-bit quantity from network into host byte order
<code>ntohs(3SOCKET)</code>	16-bit quantity from network into host byte order

The byte-swapping routines are provided because the operating system expects addresses to be supplied in network order. On some architectures, the host byte ordering is different from network byte order, so programs must sometimes byte-swap values. Routines that return network addresses do so in network order. Byte-swapping problems occur only when interpreting network addresses. For example, the following code formats a TCP or UDP port:

```
printf("port number %d\n", ntohs(sp->s_port));
```

On machines that do not need these routines, the routines are defined as null macros.

Client-Server Programs

The most common form of distributed application is the client-server model. In this scheme, client processes request services from a server process.

An alternate scheme is a service server that can eliminate dormant server processes. An example is `inetd(1M)`, the Internet service daemon. `inetd(1M)` listens at a variety of ports, determined at startup by reading a configuration file. When a connection is requested on an `inetd(1M)` serviced port, `inetd(1M)` spawns the appropriate server to serve the client. Clients are unaware that an intermediary has played any part in the connection. `inetd(1M)` is described in more detail in “`inetd Daemon`” on page 143.

Sockets and Servers

Most servers are accessed at well-known Internet port numbers or UNIX family names. The service `rlogin` is an example of a well-known UNIX family name. The main loop of a remote login server is shown in [Example 7-7](#).

The server dissociates from the controlling terminal of its invoker unless the server is operating in `DEBUG` mode.

```
(void) close(0);
(void) close(1);
(void) close(2);
(void) open("/", O_RDONLY);
(void) dup2(0, 1);
(void) dup2(0, 2);
setsid();
```

Dissociating prevents the server from receiving signals from the process group of the controlling terminal. After a server has dissociated from the controlling terminal, the server cannot send reports of errors to the terminal. The dissociated server must log errors with [syslog\(3C\)](#).

The server gets its service definition by calling [getaddrinfo\(3SOCKET\)](#).

```
bzero(&hints, sizeof (hints));
hints.ai_flags = AI_ALL|AI_ADDRCONFIG;
hints.ai_socktype = SOCK_STREAM;
error = getaddrinfo(NULL, "rlogin", &hints, &api);
```

The result, which is returned in `api`, contains the Internet port at which the program listens for service requests. Some standard port numbers are defined in `/usr/include/netinet/in.h`.

The server then creates a socket, and listens for service requests. The [bind\(3SOCKET\)](#) routine ensures that the server listens at the expected location. Because the remote login server listens at a restricted port number, the server runs as superuser. The main body of the server is the following loop.

EXAMPLE 7-7 Server Main Loop

```
/* Wait for a connection request. */
for (;;) {
    faddrlen = sizeof (faddr);
    new_sock = accept(sock, (struct sockaddr *)api->ai_addr,
                    api->ai_addrlen);
    if (new_sock == -1) {
        if (errno != EINTR && errno != ECONNABORTED) {
            perror("rlogind: accept");
        }
        continue;
    }
    if (fork() == 0) {
        close (sock);
        doit (new_sock, &faddr);
    }
    close (new_sock);
}
/*NOTREACHED*/
```

[accept\(3SOCKET\)](#) blocks messages until a client requests service. Furthermore, [accept\(3SOCKET\)](#) returns a failure indication if `accept` is interrupted by a signal, such as `SIGCHLD`. The return value from [accept\(3SOCKET\)](#) is checked, and an error is logged with [syslog\(3C\)](#), if an error occurs.

The server then forks a child process, and invokes the main body of the remote login protocol processing. The socket used by the parent to queue connection requests is closed in the child. The socket created by `accept(3SOCKET)` is closed in the parent. The address of the client is passed to the server application's `doit()` routine, which authenticates the client.

Sockets and Clients

This section describes the steps taken by a client process. As in the server, the first step is to locate the service definition for a remote login.

```
bzero(&hints, sizeof (hints));
hints.ai_flags = AI_ALL|AI_ADDRCONFIG;
hints.ai_socktype = SOCK_STREAM;

error = getaddrinfo(hostname, servicename, &hints, &res);
if (error != 0) {
    (void) fprintf(stderr, "getaddrinfo: %s for host %s service %s\n",
        gai_strerror(error), hostname, servicename);
    return (-1);
}
```

`getaddrinfo(3SOCKET)` returns the head of a list of addresses in `res`. The desired address is found by creating a socket and trying to connect to each address returned in the list until one works.

```
for (aip = res; aip != NULL; aip = aip->ai_next) {
    /*
     * Open socket. The address type depends on what
     * getaddrinfo() gave us.
     */
    sock = socket(aip->ai_family, aip->ai_socktype,
        aip->ai_protocol);
    if (sock == -1) {
        perror("socket");
        freeaddrinfo(res);
        return (-1);
    }

    /* Connect to the host. */
    if (connect(sock, aip->ai_addr, aip->ai_addrlen) == -1) {
        perror("connect");
        (void) close(sock);
        sock = -1;
        continue;
    }
    break;
}
```

The socket has been created and has been connected to the desired service. The `connect(3SOCKET)` routine implicitly binds `sock`, because `sock` is unbound.

Connectionless Servers

Some services use datagram sockets. The `rwho(1)` service provides status information on hosts that are connected to a local area network. Avoid running `in.rwhod(1M)` because `in.rwho` causes heavy network traffic. The `rwho` service broadcasts information to all hosts connected to a particular network. The `rwho` service is an example of datagram socket use.

A user on a host that is running the `rwho(1)` server can get the current status of another host with `ruptime(1)`. Typical output is illustrated in the following example.

EXAMPLE 7-8 Output of `ruptime(1)` Program

```
itchy up 9:45, 5 users, load 1.15, 1.39, 1.31
scratchy up 2+12:04, 8 users, load 4.67, 5.13, 4.59
click up 10:10, 0 users, load 0.27, 0.15, 0.14
clack up 2+06:28, 9 users, load 1.04, 1.20, 1.65
ezekiel up 25+09:48, 0 users, load 1.49, 1.43, 1.41
dandy 5+00:05, 0 users, load 1.51, 1.54, 1.56
peninsula down 0:24
wood down 17:04
carpediem down 16:09
chances up 2+15:57, 3 users, load 1.52, 1.81, 1.86
```

Status information is periodically broadcast by the `rwho(1)` server processes on each host. The server process also receives the status information. The server also updates a database. This database is interpreted for the status of each host. Servers operate autonomously, coupled only by the local network and its broadcast capabilities.

Use of broadcast is fairly inefficient because broadcast generates a lot of net traffic. Unless the service is used widely and frequently, the expense of periodic broadcasts outweighs the simplicity.

The following example shows a simplified version of the `rwho(1)` server. The sample code receives status information broadcast by other hosts on the network and supplies the status of the host on which the sample code is running. The first task is done in the main loop of the program: Packets received at the `rwho(1)` port are checked to be sure they were sent by another `rwho(1)` server process and are stamped with the arrival time. The packets then update a file with the status of the host. When a host has not been heard from for an extended time, the database routines assume the host is down and logs this information. Because a server might be down while a host is up, this application is prone to error.

EXAMPLE 7-9 `rwho(1)` Server

```
main()
{
    ...
    sp = getservbyname("who", "udp");
    net = getnetbyname("localnet");
    sin.sin6_addr = inet_makeaddr(net->n_net, in6addr_any);
    sin.sin6_port = sp->s_port;
```

EXAMPLE 7-9 rwho(1) Server (Continued)

```

...
s = socket(AF_INET6, SOCK_DGRAM, 0);
...
on = 1;
if (setsockopt(s, SOL_SOCKET, SO_BROADCAST, &on, sizeof on)
    == -1) {
    syslog(LOG_ERR, "setsockopt SO_BROADCAST: %m");
    exit(1);
}
bind(s, (struct sockaddr *) &sin, sizeof sin);
...
signal(SIGALRM, onalrm);
onalrm();
while(1) {
    struct whod wd;
    int cc, whod, len = sizeof from;
    cc = recvfrom(s, (char *) &wd, sizeof(struct whod), 0,
        (struct sockaddr *) &from, &len);
    if (cc <= 0) {
        if (cc == -1 && errno != EINTR)
            syslog(LOG_ERR, "rwhod: recv: %m");
        continue;
    }
    if (from.sin6_port != sp->s_port) {
        syslog(LOG_ERR, "rwhod: %d: bad from port",
            ntohs(from.sin6_port));
        continue;
    }
    ...
    if (!verify( wd.wd_hostname)) {
        syslog(LOG_ERR, "rwhod: bad host name from %x",
            ntohl(from.sin6_addr.s6_addr));
        continue;
    }
    (void) snprintf(path, sizeof(PATH), "%s/whod.%s", RWHODIR, wd.wd_hostname);
    whod = open(path, O_WRONLY|O_CREAT|O_TRUNC|O_NOFOLLOW, 0666);
    ...
    (void) time(&wd.wd_recvtime);
    (void) write(whod, (char *) &wd, cc);
    (void) close(whod);
}
exit(0);
}

```

The second server task is to supply the status of its host. This requires periodically acquiring system status information, packaging it in a message, and broadcasting it on the local network for other `rwho(1)` servers to hear. This task is run by a timer. This task is triggered by a signal.

Status information is broadcast on the local network. For networks that do not support broadcast, use multicast.

Advanced Socket Topics

For most programmers, the mechanisms already described are enough to build distributed applications. This section describes additional features.

Out-of-Band Data

The stream socket abstraction includes out-of-band data. Out-of-band data is a logically independent transmission channel between a pair of connected stream sockets. Out-of-band data is delivered independent of normal data. The out-of-band data facilities must support the reliable delivery of at least one out-of-band message at a time. This message can contain at least one byte of data. At least one message can be pending delivery at any time.

With in-band signaling, urgent data is delivered in sequence with normal data, and the message is extracted from the normal data stream. The extracted message is stored separately. Users can choose between receiving the urgent data in order and receiving the data out of sequence, without having to buffer the intervening data.

Using `MSG_PEEK`, you can peek at out-of-band data. If the socket has a process group, a `SIGURG` signal is generated when the protocol is notified of its existence. A process can set the process group or process ID to deliver `SIGURG` to with the appropriate `fcntl(2)` call, as described in “[Interrupt-Driven Socket I/O](#)” on page 139 for `SIGIO`. If multiple sockets have out-of-band data waiting for delivery, a `select(3C)` call for exceptional conditions can determine which sockets have such data pending.

A logical mark is placed in the data stream at the point at which the out-of-band data was sent. The remote login and remote shell applications use this facility to propagate signals between client and server processes. When a signal is received, all data up to the mark in the data stream is discarded.

To send an out-of-band message, apply the `MSG_OOB` flag to `send(3SOCKET)` or `sendto(3SOCKET)`. To receive out-of-band data, specify `MSG_OOB` to `recvfrom(3SOCKET)` or `recv(3SOCKET)`. If out-of-band data is taken in line the `MSG_OOB` flag is not needed. The `SIOCATMARK ioctl(2)` indicates whether the read pointer currently points at the mark in the data stream:

```
int yes;
ioctl(s, SIOCATMARK, &yes);
```

If `yes` is 1 on return, the next read returns data after the mark. Otherwise, assuming out-of-band data has arrived, the next read provides data sent by the client before sending the out-of-band signal. The routine in the remote login process that flushes output on receipt of an interrupt or quit signal is shown in the following example. This code reads the normal data up to the mark to discard the normal data, then reads the out-of-band byte.

A process can also read or peek at the out-of-band data without first reading up to the mark. Accessing this data when the underlying protocol delivers the urgent data in-band with the normal data, and sends notification of its presence only ahead of time, is more difficult. An example of this type of protocol is TCP, the protocol used to provide socket streams in the Internet family. With such protocols, the out-of-band byte might not yet have arrived when `recv(3SOCKET)` is called with the `MSG_OOB` flag. In that case, the call returns the error of `EWOULDBLOCK`. Also, the amount of in-band data in the input buffer might cause normal flow control to prevent the peer from sending the urgent data until the buffer is cleared. The process must then read enough of the queued data to clear the input buffer before the peer can send the urgent data.

EXAMPLE 7-10 Flushing Terminal I/O on Receipt of Out-of-Band Data

```
#include <sys/ioctl.h>
#include <sys/file.h>
...
oob()
{
    int out = FWRITE;
    char waste[BUFSIZ];
    int mark = 0;

    /* flush local terminal output */
    ioctl(1, TIOCFDUSH, (char *) &out);
    while(1) {
        if (ioctl(rem, SIOCATMARK, &mark) == -1) {
            perror("ioctl");
            break;
        }
        if (mark)
            break;
        (void) read(rem, waste, sizeof waste);
    }
    if (recv(rem, &mark, 1, MSG_OOB) == -1) {
        perror("recv");
        ...
    }
    ...
}
```

A facility to retain the position of urgent in-line data in the socket stream is available as a socket-level option, `SO_OOBINLINE`. See `getsockopt(3SOCKET)` for usage. With this socket-level option, the position of urgent data remains. However, the urgent data immediately following the mark in the normal data stream is returned without the `MSG_OOB` flag. Reception of multiple urgent indications moves the mark, but does not lose any out-of-band data.

Nonblocking Sockets

Some applications require sockets that do not block. For example, a server would return an error code, not executing a request that cannot complete immediately. This error could cause

the process to be suspended, awaiting completion. After creating and connecting a socket, issuing a `fcntl(2)` call, as shown in the following example, makes the socket nonblocking.

EXAMPLE 7-11 Set Nonblocking Socket

```
#include <fcntl.h>
#include <sys/file.h>
...
int fileflags;
int s;
...
s = socket(AF_INET6, SOCK_STREAM, 0);
...
if (fileflags = fcntl(s, F_GETFL, 0) == -1)
    perror("fcntl F_GETFL");
    exit(1);
}
if (fcntl(s, F_SETFL, fileflags | FNDELAY) == -1)
    perror("fcntl F_SETFL, FNDELAY");
    exit(1);
}
```

When performing I/O on a nonblocking socket, check for the error `EWOULDBLOCK` in `errno.h`, which occurs when an operation would normally block. `accept(3SOCKET)`, `connect(3SOCKET)`, `send(3SOCKET)`, `recv(3SOCKET)`, `read(2)`, and `write(2)` can all return `EWOULDBLOCK`. If an operation such as a `send(3SOCKET)` cannot be done in its entirety but partial writes work, as when using a stream socket, all available data is processed. The return value is the amount of data actually sent.

Asynchronous Socket I/O

Asynchronous communication between processes is required in applications that simultaneously handle multiple requests. Asynchronous sockets must be of the `SOCK_STREAM` type. To make a socket asynchronous, you issue a `fcntl(2)` call, as shown in the following example.

EXAMPLE 7-12 Making a Socket Asynchronous

```
#include <fcntl.h>
#include <sys/file.h>
...
int fileflags;
int s;
...
s = socket(AF_INET6, SOCK_STREAM, 0);
...
if (fileflags = fcntl(s, F_GETFL ) == -1)
    perror("fcntl F_GETFL");
    exit(1);
}
if (fcntl(s, F_SETFL, fileflags | FNDELAY | FASYNC) == -1)
```

EXAMPLE 7-12 Making a Socket Asynchronous (Continued)

```

    perror("fcntl F_SETFL, FNDELAY | FASYNC");
    exit(1);
}

```

After sockets are initialized, connected, and made nonblocking and asynchronous, communication is similar to reading and writing a file asynchronously. Initiate a data transfer by using `send(3SOCKET)`, `write(2)`, `recv(3SOCKET)`, or `read(2)`. A signal-driven I/O routine completes a data transfer, as described in the next section.

Interrupt-Driven Socket I/O

The SIGIO signal notifies a process when a socket, or any file descriptor, has finished a data transfer. The steps in using SIGIO are as follows:

1. Set up a SIGIO signal handler with the `signal(3C)` or `sigvec` calls.
2. Use `fcntl(2)` to set the process ID or process group ID to route the signal to its own process ID or process group ID. The default process group of a socket is group 0.
3. Convert the socket to asynchronous, as shown in “Asynchronous Socket I/O” on page 138.

The following sample code enables receipt of information on pending requests as the requests occur for a socket by a given process. With the addition of a handler for SIGURG, this code can also be used to prepare for receipt of SIGURG signals.

EXAMPLE 7-13 Asynchronous Notification of I/O Requests

```

#include <fcntl.h>
#include <sys/file.h>
...
signal(SIGIO, io_handler);
/* Set the process receiving SIGIO/SIGURG signals to us. */
if (fcntl(s, F_SETOWN, getpid()) < 0) {
    perror("fcntl F_SETOWN");
    exit(1);
}

```

Signals and Process Group ID

For SIGURG and SIGIO, each socket has a process number and a process group ID. These values are initialized to zero, but can be redefined at a later time with the `F_SETOWN fcntl(2)` command, as in the previous example. A positive third argument to `fcntl(2)` sets the socket's process ID. A negative third argument to `fcntl(2)` sets the socket's process group ID. The only allowed recipient of SIGURG and SIGIO signals is the calling process. A similar `fcntl(2)`, `F_GETOWN`, returns the process number of a socket.

You can also enable reception of SIGURG and SIGIO by using `ioctl(2)` to assign the socket to the user's process group.

```
/* oobdata is the out-of-band data handling routine */
sigset(SIGURG, oobdata);
int pid = -getpid();
if (ioctl(client, SIOCSGRP, (char *) &pid) < 0) {
    perror("ioctl: SIOCSGRP");
}
```

Selecting Specific Protocols

If the third argument of the `socket(3SOCKET)` call is 0, `socket(3SOCKET)` selects a default protocol to use with the returned socket of the type requested. The default protocol is usually correct, and alternate choices are not usually available. When using raw sockets to communicate directly with lower-level protocols or lower-level hardware interfaces, set up de-multiplexing with the protocol argument.

Using raw sockets in the Internet family to implement a new protocol on IP ensures that the socket only receives packets for the specified protocol. To obtain a particular protocol, determine the protocol number as defined in the protocol family. For the Internet family, use one of the library routines that are discussed in “Standard Routines” on page 127, such as `getprotobyname(3SOCKET)`.

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netdb.h>
...
pp = getprotobyname("newtcp");
s = socket(AF_INET6, SOCK_STREAM, pp->p_proto);
```

Using `getprotobyname` results in a socket `s` by using a stream-based connection, but with a protocol type of `newtcp` instead of the default `tcp`.

Address Binding

For addressing, TCP and UDP use a 4-tuple of:

- Local IP address
- Local port number
- Foreign IP address
- Foreign port number

TCP requires these 4-tuples to be unique. UDP does not. User programs do not always know proper values to use for the local address and local port, because a host can reside on multiple networks. The set of allocated port numbers is not directly accessible to a user. To avoid these

problems, leave parts of the address unspecified and let the system assign the parts appropriately when needed. Various portions of these tuples can be specified by various parts of the sockets API:

<code>bind(3SOCKET)</code>	Local address or local port or both
<code>connect(3SOCKET)</code>	Foreign address and foreign port

A call to `accept(3SOCKET)` retrieves connection information from a foreign client. This causes the local address and port to be specified to the system even though the caller of `accept(3SOCKET)` did not specify anything. The foreign address and foreign port are returned.

A call to `listen(3SOCKET)` can cause a local port to be chosen. If no explicit `bind(3SOCKET)` has been done to assign local information, `listen(3SOCKET)` assigns an ephemeral port number.

A service that resides at a particular port can `bind(3SOCKET)` to that port. Such a service can leave the local address unspecified if the service does not require local address information. The local address is set to `in6addr_any`, a variable with a constant value in `<netinet/in.h>`. If the local port does not need to be fixed, a call to `listen(3SOCKET)` causes a port to be chosen. Specifying an address of `in6addr_any` or a port number of 0 is known as *wildcarding*. For `AF_INET`, `INADDR_ANY` is used in place of `in6addr_any`.

The wildcard address simplifies local address binding in the Internet family. The following sample code binds a specific port number that was returned by a call to `getaddrinfo(3SOCKET)` to a socket and leaves the local address unspecified:

```
#include <sys/types.h>
#include <netinet/in.h>
...
struct addrinfo *aip;
...
if (bind(sock, aip->ai_addr, aip->ai_addrlen) == -1) {
    perror("bind");
    (void) close(sock);
    return (-1);
}
```

Each network interface on a host typically has a unique IP address. Sockets with wildcard local addresses can receive messages that are directed to the specified port number. Messages that are sent to any of the possible addresses that are assigned to a host are also received by sockets with wildcard local addresses. To allow only hosts on a specific network to connect to the server, a server binds the address of the interface on the appropriate network.

Similarly, a local port number can be left unspecified, in which case the system selects a port number. For example, to bind a specific local address to a socket, but to leave the local port number unspecified, you could use `bind` as follows:

```
bzero (&sin, sizeof (sin));
(void) inet_pton (AF_INET6, "::ffff:127.0.0.1", sin.sin6_addr.s6_addr);
sin.sin6_family = AF_INET6;
sin.sin6_port = htons(0);
bind(s, (struct sockaddr *) &sin, sizeof sin);
```

The system uses two criteria to select the local port number:

- Internet port numbers less than 1024 (IPPORT_RESERVED) are reserved for privileged users. Nonprivileged users can use any Internet port number that is greater than 1024. The largest Internet port number is 65535.
- The port number is not currently bound to some other socket.

The port number and IP address of the client are found through either [accept\(3SOCKET\)](#) or [getpeername\(3SOCKET\)](#).

In certain cases, the algorithm used by the system to select port numbers is unsuitable for an application due to the two-step creation process for associations. For example, the Internet file transfer protocol specifies that data connections must always originate from the same local port. However, duplicate associations are avoided by connecting to different foreign ports. In this situation, the system would disallow binding the same local address and local port number to a socket if a previous data connection's socket still existed.

To override the default port selection algorithm, you must perform an option call before address binding:

```
int on = 1;
...
setsockopt(s, SOL_SOCKET, SO_REUSEADDR, &on, sizeof on);
bind(s, (struct sockaddr *) &sin, sizeof sin);
```

With this call, local addresses already in use can be bound. This binding does not violate the uniqueness requirement. The system still verifies at connect time that any other sockets with the same local address and local port do not have the same foreign address and foreign port. If the association already exists, the error EADDRINUSE is returned.

Socket Options

You can set and get several options on sockets through [setsockopt\(3SOCKET\)](#) and [getsockopt\(3SOCKET\)](#). For example, you can change the send or receive buffer space. The general forms of the calls are in the following list:

```
setsockopt(s, level, optname, optval, optlen);
```

and

```
getsockopt(s, level, optname, optval, optlen);
```

The operating system can adjust the values appropriately at any time.

The arguments of `setsockopt(3SOCKET)` and `getsockopt(3SOCKET)` calls are in the following list:

<i>s</i>	Socket on which the option is to be applied
<i>level</i>	Specifies the protocol level, such as socket level, indicated by the symbolic constant <code>SOL_SOCKET</code> in <code>sys/socket.h</code>
<i>optname</i>	Symbolic constant defined in <code>sys/socket.h</code> that specifies the option
<i>optval</i>	Points to the value of the option
<i>optlen</i>	Points to the length of the value of the option

For `getsockopt(3SOCKET)`, *optlen* is a value-result argument. This argument is initially set to the size of the storage area pointed to by *optval*. On return, the argument's value is set to the length of storage used.

When a program needs to determine an existing socket's type, the program should invoke `inetd(1M)` by using the `SO_TYPE` socket option and the `getsockopt(3SOCKET)` call:

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

int type, size;

size = sizeof (int);
if (getsockopt(s, SOL_SOCKET, SO_TYPE, (char *) &type, &size) < 0) {
    ...
}
```

After `getsockopt(3SOCKET)`, *type* is set to the value of the socket type, as defined in `sys/socket.h`. For a datagram socket, *type* would be `SOCK_DGRAM`.

inetd Daemon

The `inetd(1M)` daemon is invoked at startup time and is now configured via `smf(5)`. The configuration was previously performed by `/etc/inet/inetd.conf` file.

Use `inetconv(1M)` to convert the configuration file content into SMF format services, and then manage these services using `inetadm(1M)` and `svcadm(1M)`. See the `inetd(1M)` man page for details.

The `inetd(1M)` daemon polls each socket, waiting for a connection request to the service corresponding to that socket. For `SOCK_STREAM` type sockets, `inetd(1M)` accepts

(`accept(3SOCKET)`) on the listening socket, forks (`fork(2)`), duplicates (`dup(2)`) the new socket to file descriptors 0 and 1 (`stdin` and `stdout`), closes other open file descriptors, and executes (`exec(2)`) the appropriate server.

The primary benefit of using `inetd(1M)` is that services not in use do not consume machine resources. A secondary benefit is that `inetd(1M)` does most of the work to establish a connection. The server started by `inetd(1M)` has the socket connected to its client on file descriptors 0 and 1. The server can immediately read, write, send, or receive. Servers can use buffered I/O as provided by the `stdio` conventions, as long as the servers use `fflush(3C)` when appropriate.

The `getpeername(3SOCKET)` routine returns the address of the peer (process) connected to a socket. This routine is useful in servers started by `inetd(1M)`. For example, you could use this routine to log the Internet address such as `fec0::56:a00:20ff:fe7d:3dd2`, which is conventional for representing the IPv6 address of a client. An `inetd(1M)` server could use the following sample code:

```
struct sockaddr_storage name;
int namelen = sizeof (name);
char abuf[INET6_ADDRSTRLEN];
struct in6_addr addr6;
struct in_addr addr;

if (getpeername(fd, (struct sockaddr *) &name, &namelen) == -1) {
    perror("getpeername");
    exit(1);
} else {
    addr = ((struct sockaddr_in *)&name)->sin_addr;
    addr6 = ((struct sockaddr_in6 *)&name)->sin6_addr;
    if (name.ss_family == AF_INET) {
        (void) inet_ntop(AF_INET, &addr, abuf, sizeof (abuf));
    } else if (name.ss_family == AF_INET6 &&
               IN6_IS_ADDR_V4MAPPED(&addr6)) {
        /* this is a IPv4-mapped IPv6 address */
        IN6_MAPPED_TO_IN(&addr6, &addr);
        (void) inet_ntop(AF_INET, &addr, abuf, sizeof (abuf));
    } else if (name.ss_family == AF_INET6) {
        (void) inet_ntop(AF_INET6, &addr6, abuf, sizeof (abuf));
    }
    syslog("Connection from %s\n", abuf);
}
```

Broadcasting and Determining Network Configuration

Broadcasting is not supported in IPv6. Broadcasting is supported only in IPv4.

Messages sent by datagram sockets can be broadcast to reach all of the hosts on an attached network. The network must support broadcast because the system provides no simulation of broadcast in software. Broadcast messages can place a high load on a network because broadcast messages force every host on the network to service the broadcast messages. Broadcasting is usually used for either of two reasons:

- To find a resource on a local network without having its address
- For functions that require information to be sent to all accessible neighbors

To send a broadcast message, create an Internet datagram socket:

```
s = socket(AF_INET, SOCK_DGRAM, 0);
```

Bind a port number to the socket:

```
sin.sin_family = AF_INET;
sin.sin_addr.s_addr = htonl(INADDR_ANY);
sin.sin_port = htons(MYPORT);
bind(s, (struct sockaddr *) &sin, sizeof sin);
```

The datagram can be broadcast on only one network by sending to the network's broadcast address. A datagram can also be broadcast on all attached networks by sending to the special address `INADDR_BROADCAST`, which is defined in `netinet/in.h`.

The system provides a mechanism to determine a number of pieces of information about the network interfaces on the system. This information includes the IP address and broadcast address. The `SIOCGIFCONF` `ioctl(2)` call returns the interface configuration of a host in a single `ifconf` structure. This structure contains an array of `ifreq` structures. Every address family supported by every network interface to which the host is connected has its own `ifreq` structure.

The following example shows the `ifreq` structures defined in `net/if.h`.

EXAMPLE 7-14 net/if.h Header File

```
struct ifreq {
    #define IFNAMSIZ 16
    char ifr_name[IFNAMSIZ]; /* if name, e.g., "en0" */
    union {
        struct sockaddr ifru_addr;
        struct sockaddr ifru_dstaddr;
        char ifru_ename[IFNAMSIZ]; /* other if name */
        struct sockaddr ifru_broadaddr;
        short ifru_flags;
        int ifru_metric;
        char ifru_data[1]; /* interface dependent data */
        char ifru_enaddr[6];
    } ifr_ifru;
    #define ifr_addr ifr_ifru.ifru_addr
    #define ifr_dstaddr ifr_ifru.ifru_dstaddr
    #define ifr_ename ifr_ifru.ifru_ename
```

EXAMPLE 7-14 net/if.h Header File (Continued)

```

#define ifr_broadaddr ifr_ifru.ifru_broadaddr
#define ifr_flags ifr_ifru.ifru_flags
#define ifr_metric ifr_ifru.ifru_metric
#define ifr_data ifr_ifru.ifru_data
#define ifr_enaddr ifr_ifru.ifru_enaddr
};

```

The call that obtains the interface configuration is:

```

/*
 * Do SIOCGIFNUM ioctl to find the number of interfaces
 *
 * Allocate space for number of interfaces found
 *
 * Do SIOCGIFCONF with allocated buffer
 *
 */
if (ioctl(s, SIOCGIFNUM, (char *)&numifs) == -1) {
    numifs = MAXIFS;
}
bufsize = numifs * sizeof(struct ifreq);
reqbuf = (struct ifreq *)malloc(bufsize);
if (reqbuf == NULL) {
    fprintf(stderr, "out of memory\n");
    exit(1);
}
ifc.ifc_buf = (caddr_t)&reqbuf[0];
ifc.ifc_len = bufsize;
if (ioctl(s, SIOCGIFCONF, (char *)&ifc) == -1) {
    perror("ioctl(SIOCGIFCONF)");
    exit(1);
}

```

After this call, *buf* contains an array of *ifreq* structures. Every network to which the host connects has an associated *ifreq* structure. The sort order these structures appear in is:

- Alphabetical by interface name
- Numerical by supported address families

The value of *ifc.ifc_len* is set to the number of bytes used by the *ifreq* structures.

Each structure has a set of interface flags that indicate whether the corresponding network is up or down, point-to-point or broadcast, and so on. The following example shows `ioctl(2)` returning the `SIOCGIFFLAGS` flags for an interface specified by an *ifreq* structure.

EXAMPLE 7-15 Obtaining Interface Flags

```

struct ifreq *ifr;
ifr = ifc.ifc_req;
for (n = ifc.ifc_len/sizeof (struct ifreq); --n >= 0; ifr++) {
    /*

```

EXAMPLE 7-15 Obtaining Interface Flags (Continued)

```

    * Be careful not to use an interface devoted to an address
    * family other than those intended.
    */
    if (ifr->ifr_addr.sa_family != AF_INET)
        continue;
    if (ioctl(s, SIOCGIFFLAGS, (char *) ifr) < 0) {
        ...
    }
    /* Skip boring cases */
    if ((ifr->ifr_flags & IFF_UP) == 0 ||
        (ifr->ifr_flags & IFF_LOOPBACK) ||
        (ifr->ifr_flags & (IFF_BROADCAST | IFF_POINTOPOINT)) == 0)
        continue;
}

```

The following example uses the `SIOCGIFBRDADDR` `ioctl(2)` command to obtain the broadcast address of an interface.

EXAMPLE 7-16 Broadcast Address of an Interface

```

if (ioctl(s, SIOCGIFBRDADDR, (char *) ifr) < 0) {
    ...
}
memcpy((char *) &dst, (char *) &ifr->ifr_broadaddr,
       sizeof ifr->ifr_broadaddr);

```

You can also use `SIOCGIFBRDADDR` `ioctl(2)` to get the destination address of a point-to-point interface.

After the interface broadcast address is obtained, transmit the broadcast datagram with `sendto(3SOCKET)`:

```
sendto(s, buf, buflen, 0, (struct sockaddr *)&dst, sizeof dst);
```

Use one `sendto(3SOCKET)` for each interface to which the host is connected, if that interface supports the broadcast or point-to-point addressing.

Using Multicast

IP multicasting is supported only on `AF_INET6` and `AF_INET` sockets of type `SOCK_DGRAM` and `SOCK_RAW`. IP multicasting is only supported on subnetworks for which the interface driver supports multicasting.

Sending IPv4 Multicast Datagrams

To send a multicast datagram, specify an IP multicast address in the range 224.0.0.0 to 239.255.255.255 as the destination address in a `sendto(3SOCKET)` call.

By default, IP multicast datagrams are sent with a time-to-live (TTL) of 1. This value prevents the datagrams from being forwarded beyond a single subnetwork. The socket option `IP_MULTICAST_TTL` allows the TTL for subsequent multicast datagrams to be set to any value from 0 to 255. This ability is used to control the scope of the multicasts.

```
u_char ttl;
setsockopt(sock, IPPROTO_IP, IP_MULTICAST_TTL, &ttl, sizeof(ttl))
```

Multicast datagrams with a TTL of 0 are not transmitted on any subnet, but can be delivered locally if the sending host belongs to the destination group and if multicast loopback has not been disabled on the sending socket. Multicast datagrams with a TTL greater than one can be delivered to more than one subnet if one or more multicast routers are attached to the first-hop subnet. To provide meaningful scope control, the multicast routers support the notion of TTL thresholds. These thresholds prevent datagrams with less than a certain TTL from traversing certain subnets. The thresholds enforce the conventions for multicast datagrams with initial TTL values as follows:

0	Are restricted to the same host
1	Are restricted to the same subnet
32	Are restricted to the same site
64	Are restricted to the same region
128	Are restricted to the same continent
255	Are unrestricted in scope

Sites and regions are not strictly defined and sites can be subdivided into smaller administrative units as a local matter.

An application can choose an initial TTL other than the ones previously listed. For example, an application might perform an expanding-ring search for a network resource by sending a multicast query, first with a TTL of 0 and then with larger and larger TTLs until a reply is received.

The multicast router does not forward any multicast datagram with a destination address between 224.0.0.0 and 224.0.0.255 inclusive, regardless of its TTL. This range of addresses is reserved for the use of routing protocols and other low-level topology discovery or maintenance protocols, such as gateway discovery and group membership reporting.

Each multicast transmission is sent from a single network interface, even if the host has more than one multicast-capable interface. If the host is also a multicast router and the TTL is greater than 1, a multicast can be *forwarded* to interfaces other than the originating interface. A socket option is available to override the default for subsequent transmissions from a given socket:

```
struct in_addr addr;
setsockopt(sock, IPPROTO_IP, IP_MULTICAST_IF, &addr, sizeof(addr))
```

where `addr` is the local IP address of the desired outgoing interface. Revert to the default interface by specifying the address `INADDR_ANY`. The local IP address of an interface is obtained with the `SIOCGIFCONF` `ioctl`. To determine if an interface supports multicasting, fetch the interface flags with the `SIOCGIFFLAGS` `ioctl` and test if the `IFF_MULTICAST` flag is set. This option is intended primarily for multicast routers and other system services specifically concerned with Internet topology.

If a multicast datagram is sent to a group to which the sending host itself belongs, a copy of the datagram is, by default, looped back by the IP layer for local delivery. Another socket option gives the sender explicit control over whether subsequent datagrams are looped back:

```
u_char loop;
setsockopt(sock, IPPROTO_IP, IP_MULTICAST_LOOP, &loop, sizeof(loop))
```

where `loop` is 0 to disable loopback and 1 to enable loopback. This option provides a performance benefit for applications that have only one instance on a single host by eliminating the overhead of receiving their own transmissions. Applications that can have more than one instance on a single host, or for which the sender does not belong to the destination group, should not use this option.

If the sending host belongs to the destination group on another interface, a multicast datagram sent with an initial TTL greater than 1 can be delivered to the sending host on the other interface. The loopback control option has no effect on such delivery.

Receiving IPv4 Multicast Datagrams

Before a host can receive IP multicast datagrams, the host must become a member of one or more IP multicast groups. A process can ask the host to join a multicast group by using the following socket option:

```
struct ip_mreq mreq;
setsockopt(sock, IPPROTO_IP, IP_ADD_MEMBERSHIP, &mreq, sizeof(mreq))
```

where `mreq` is the structure:

```
struct ip_mreq {
    struct in_addr imr_multiaddr;    /* multicast group to join */
```

```
    struct in_addr imr_interface;    /* interface to join on */  
}
```

Each membership is associated with a single interface. You can join the same group on more than one interface. Specify the `imr_interface` address as `INADDR_ANY` to choose the default multicast interface. You can also specify one of the host's local addresses to choose a particular multicast-capable interface.

To drop a membership, use:

```
struct ip_mreq mreq;  
setsockopt(sock, IPPROTO_IP, IP_DROP_MEMBERSHIP, &mreq, sizeof(mreq))
```

where `mreq` contains the same values used to add the membership. Closing a socket or killing the process that holds the socket drops the memberships associated with that socket. More than one socket can claim a membership in a particular group, and the host remains a member of that group until the last claim is dropped.

If any socket claims membership in the destination group of the datagram, the kernel IP layer accepts incoming multicast packets. A given socket's receipt of a multicast datagram depends on the socket's associated destination port and memberships, or the protocol type for raw sockets. To receive multicast datagrams sent to a particular port, bind to the local port, leaving the local address unspecified, such as `INADDR_ANY`.

More than one process can bind to the same `SOCK_DGRAM` UDP port if the `bind(3SOCKET)` is preceded by:

```
int one = 1;  
setsockopt(sock, SOL_SOCKET, SO_REUSEADDR, &one, sizeof(one))
```

In this case, every incoming multicast or broadcast UDP datagram destined for the shared port is delivered to all sockets bound to the port. For backwards compatibility reasons, this delivery does *not* apply to incoming unicast datagrams. Unicast datagrams are never delivered to more than one socket, regardless of how many sockets are bound to the datagram's destination port. `SOCK_RAW` sockets do not require the `SO_REUSEADDR` option to share a single IP protocol type.

The definitions required for the new, multicast-related socket options are found in `<netinet/in.h>`. All IP addresses are passed in network byte-order.

Sending IPv6 Multicast Datagrams

To send an IPv6 multicast datagram, specify an IP multicast address in the range `ff00::0/8` as the destination address in a `sendto(3SOCKET)` call.

By default, IP multicast datagrams are sent with a hop limit of one, which prevents the datagrams from being forwarded beyond a single subnetwork. The socket option

IPV6_MULTICAST_HOPS allows the hop limit for subsequent multicast datagrams to be set to any value from 0 to 255. This ability is used to control the scope of the multicasts:

```
uint_l;
setsockopt(sock, IPPROTO_IPV6, IPV6_MULTICAST_HOPS, &hops, sizeof(hops))
```

You cannot transmit multicast datagrams with a hop limit of zero on any subnet, but you can deliver the datagrams locally if:

- The sending host belongs to the destination group
- Multicast loopback on the sending socket is enabled

You can deliver multicast datagrams with a hop limit that is greater than one to more than one subnet if the first-hop subnet attaches to one or more multicast routers. The IPv6 multicast addresses, unlike their IPv4 counterparts, contain explicit scope information that is encoded in the first part of the address. The defined scopes are, where X is unspecified:

```
ffX1::0/16    Node-local scope, restricted to the same node
ffX2::0/16    Link-local scope
ffX5::0/16    Site-local scope
ffX8::0/16    Organization-local scope
ffXe::0/16    Global scope
```

An application can, separately from the scope of the multicast address, use different hop limit values. For example, an application might perform an expanding-ring search for a network resource by sending a multicast query, first with a hop limit of 0, and then with larger and larger hop limits, until a reply is received.

Each multicast transmission is sent from a single network interface, even if the host has more than one multicast-capable interface. If the host is also a multicast router, and the hop limit is greater than 1, a multicast can be *forwarded* to interfaces other than the originating interface. A socket option is available to override the default for subsequent transmissions from a given socket:

```
uint_t ifindex;
ifindex = if_nametoindex("hme3");
setsockopt(sock, IPPROTO_IPV6, IPV6_MULTICAST_IF, &ifindex, sizeof(ifindex))
```

where `ifindex` is the interface index for the desired outgoing interface. Revert to the default interface by specifying the value 0.

If a multicast datagram is sent to a group to which the sending host itself belongs, a copy of the datagram is, by default, looped back by the IP layer for local delivery. Another socket option gives the sender explicit control over whether to loop back subsequent datagrams:

```
uint_t loop;
setsockopt(sock, IPPROTO_IPV6, IPV6_MULTICAST_LOOP, &loop, sizeof(loop))
```

where `loop` is zero to disable loopback and one to enable loopback. This option provides a performance benefit for applications that have only one instance on a single host (such as a router or a mail demon), by eliminating the overhead of receiving their own transmissions. Applications that can have more than one instance on a single host (such as a conferencing program) or for which the sender does not belong to the destination group (such as a time querying program) should not use this option.

If the sending host belongs to the destination group on another interface, a multicast datagram sent with an initial hop limit greater than 1 can be delivered to the sending host on the other interface. The loopback control option has no effect on such delivery.

Receiving IPv6 Multicast Datagrams

Before a host can receive IP multicast datagrams, the host must become a member of one, or more IP multicast groups. A process can ask the host to join a multicast group by using the following socket option:

```
struct ipv6_mreq mreq;
setsockopt(sock, IPPROTO_IPV6, IPV6_JOIN_GROUP, &mreq, sizeof(mreq))
```

where `mreq` is the structure:

```
struct ipv6_mreq {
    struct in6_addr    ipv6mr_multiaddr;    /* IPv6 multicast addr */
    unsigned int       ipv6mr_interface;    /* interface index */
}
```

Each membership is associated with a single interface. You can join the same group on more than one interface. Specify `ipv6_interface` to be 0 to choose the default multicast interface. Specify an interface index for one of the host's interfaces to choose that multicast-capable interface.

To leave a group, use:

```
struct ipv6_mreq mreq;
setsockopt(sock, IPPROTO_IPV6, IP_LEAVE_GROUP, &mreq, sizeof(mreq))
```

where `mreq` contains the same values used to add the membership. The socket drops associated memberships when the socket is closed, or when the process that holds the socket is killed. More than one socket can claim a membership in a particular group. The host remains a member of that group until the last claim is dropped.

The kernel IP layer accepts incoming multicast packets if any socket has claimed a membership in the destination group of the datagram. Delivery of a multicast datagram to a particular socket

is determined by the destination port and the memberships associated with the socket, or by the protocol type for raw sockets. To receive multicast datagrams sent to a particular port, bind to the local port, leaving the local address unspecified, such as `INADDR_ANY`.

More than one process can bind to the same `SOCK_DGRAM` UDP port if the `bind(3SOCKET)` is preceded by:

```
int one = 1;
setsockopt(sock, SOL_SOCKET, SO_REUSEADDR, &one, sizeof(one))
```

In this case, all sockets that are bound to the port receive every incoming multicast UDP datagram destined to the shared port. For backward compatibility reasons, this delivery does *not* apply to incoming unicast datagrams. Unicast datagrams are never delivered to more than one socket, regardless of how many sockets are bound to the datagram's destination port. `SOCK_RAW` sockets do not require the `SO_REUSEADDR` option to share a single IP protocol type.

The definitions required for the new, multicast-related socket options are found in `<netinet/in.h>`. All IP addresses are passed in network byte-order.

Stream Control Transmission Protocol

Stream Control Transmission Protocol (SCTP) is a reliable transport protocol that provides services similar to the services provided by TCP. In addition, SCTP provides network-level fault tolerance. SCTP supports multihoming at either end of an association. The SCTP socket API supports a one-to-one socket style modeled after TCP. The SCTP socket API also supports a one-to-many socket style designed for use with signaling. The one-to-many socket style reduces the number of file descriptors used in a process. You must link the `libsctp` library to use SCTP function calls.

An SCTP association is set up between two endpoints. The endpoints use a four-way handshake mechanism that uses a cookie to guard against some types of denial-of-service (DoS) attacks. The endpoints can be represented by multiple IP addresses.

SCTP Stack Implementation

This section lists the details of the Oracle Solaris implementation of the IETF standard for the Stream Control Transmission Protocol (RFC 4960). The table in this section lists all of the exceptions from RFC 4960. The SCTP protocol in the Oracle Solaris operating system fully implements any sections of RFC 4960 that is not explicitly mentioned in the table.

TABLE 7-3 Oracle Solaris SCTP Implementation Exceptions from RFC 4960

RFC 4960 Section	Exceptions in the Oracle Solaris Implementation
3. SCTP Packet Format	<p>3.2 Chunk Field Descriptions: Oracle Solaris SCTP does not implement the optional ECNE and CWR.</p> <p>3.3.2: Oracle Solaris SCTP does not implement the Initiation (INIT) Optional ECN, Host Name Address, and Cookie Preserving parameters.</p> <p>3.3.3: Oracle Solaris SCTP does not implement the Initiation Acknowledgement, Optional ECN, and Host Name Address parameters.</p>
5. Association Initialization	5.1.2, Handle Address Parameters: Section (B), Optional Host Name parameter, is not implemented.
10. Interface with Upper Layer	Oracle Solaris SCTP implements the IETF TSVWG SCTP socket API draft.

Note – The Oracle Solaris 11 implementation of the TSVWG SCTP socket API is based on a version of the API draft that was published at the time when Oracle Solaris 11 was first shipped.

SCTP Socket Interfaces

When the `socket()` call creates a socket for `IPPROTO_SCTP`, it calls an SCTP-specific socket creation routine. Socket calls made on an SCTP socket call the appropriate SCTP socket routine automatically. In a one-to-one socket, each socket corresponds to one SCTP association. Create a one-to-one socket by calling this function:

```
socket(AF_INET[6], SOCK_STREAM, IPPROTO_SCTP);
```

In a one-to-many style socket, each socket handles multiple SCTP associations. Each association has an association identifier called `sctp_assoc_t`. Create a one-to-many socket by calling this function:

```
socket(AF_INET[6], SOCK_SEQPACKET, IPPROTO_SCTP);
```

sctp_bindx()

```
int sctp_bindx(int sock, void *addrs, int addrcnt, int flags);
```

The `sctp_bindx()` function manages addresses on an SCTP socket. If the `sock` parameter is an IPv4 socket, the addresses passed to the `sctp_bindx()` function must be IPv4 addresses.

If the *sock* parameter is an IPv6 socket, the addresses passed to the `sctp_bindx()` function can be either IPv4 (in IPv4-mapped address format) or IPv6 addresses. When the address that is passed to the `sctp_bindx()` function is `INADDR_ANY` or `IN6ADDR_ANY`, the socket binds to all available addresses. Bind the SCTP endpoint with the `bind(3SOCKET)`

If the *sock* parameter is an IPv4 socket, **addrs* should be an array of `sockaddr_in` structures containing IPv4 addresses. If *sock* is an IPv6 socket, **addrs* should be an array of `sockaddr_in6` structures containing IPv6 or IPv4-mapped IPv6 addresses. The *addrcnt* is the number of array elements in *addrs*. The family of the address type is used with *addrcnt* to determine the size of the array.

If the addresses are IPv6 addresses, they are contained in `sockaddr_in6` structures. The address type's family distinguishes the address length. The caller specifies the number of addresses in the array with the *addrcnt* parameter.

The `sctp_bindx()` function returns 0 on success. The `sctp_bindx()` function returns -1 on failure and sets the value of *errno* to the appropriate error code.

If the same port is not given for each socket address, the `sctp_bindx()` function fails and sets the value of *errno* to `EINVAL`.

The *flags* parameter is formed from performing the bitwise OR operation on zero or more of the following currently defined flags:

- `SCTP_BINDX_ADD_ADDR`
- `SCTP_BINDX_REM_ADDR`

`SCTP_BINDX_ADD_ADDR` directs SCTP to add the given addresses to the association.

`SCTP_BINDX_REM_ADDR` directs SCTP to remove the given addresses from the association. The two flags are mutually exclusive. If both are given, the `sctp_bindx()` fails and sets the value of *errno* to `EINVAL`.

The caller should add or remove addresses one at a time. If an error occurs, and a list of addresses has been used, it is not possible for the caller to find the address that caused the error. Adding or removing addresses one at a time helps the caller resolve this issue.

A caller may not remove all addresses from an association. The `sctp_bindx()` function rejects such an attempt by failing and setting the value of *errno* to `EINVAL`. An application can use `sctp_bindx(SCTP_BINDX_ADD_ADDR)` to associate additional addresses with an endpoint after calling the `bind()` function. An application can use `sctp_bindx(SCTP_BINDX_REM_ADDR)` to remove addresses associated with a listening socket. After using `sctp_bindx(SCTP_BINDX_REM_ADDR)` to remove addresses, accepting new associations will not reassociate the removed address. If the endpoint supports dynamic address, using `SCTP_BINDX_REM_ADDR` or `SCTP_BINDX_ADD_ADDR` sends a message to the peer to change the peer's address lists. Adding and removing addresses from a connected association is optional functionality. Implementations that do not support this functionality return `EOPNOTSUPP`.

If the address family is not `AF_INET` or `AF_INET6`, the `sctp_bindx()` function fails and returns `EAFNOSUPPORT`. If the file descriptor passed to the `sctp_bindx()` in the `sock` parameter is invalid, the `sctp_bindx()` function fails and returns `EBADF`.

sctp_opt_info()

```
int sctp_opt_info(int sock, sctp_assoc_id_t id, int opt, void *arg, socklen_t *len);
```

The `sctp_opt_info()` function returns the SCTP level options that are associated with the socket described in the `sock` parameter. If the socket is a one-to-many style SCTP socket the value of the `id` parameter refers to a specific association. The `id` parameter is ignored for one-to-one style SCTP sockets. The value of the `opt` parameter specifies the SCTP socket option to get. The value of the `arg` parameter is an option-specific structure buffer that is allocated by the calling program. The value of the `*len` parameter is the length of the option.

The `opt` parameter can have the following values:

- `SCTP_RTOINFO`
- `SCTP_ASSOCINFO`
- `SCTP_DEFAULT_SEND_PARAM`
- `SCTP_PEER_ADDR_PARAMS`
- `SCTP_STATUS`
- `SCTP_INITMSG`
- `SCTP_NODELAY`
- `SCTP_AUTOCLOSE`
- `SCTP_PRIMARY_ADDR`
- `SCTP_GET_PEER_ADDR_INFO`
- `SCTP_EVENT`
- `SCTP_DELAYED_SACK`
- `SCTP_PARTIAL_DELIVERY_POINT`
- `SCTP_FRAGMENT_INTERLEAVE`
- `SCTP_MAX_BURST`
- `SCTP_CONTEXT`
- `SCTP_EXPLICIT_EOR`
- `SCTP_REUSE_PORT`
- `SCTP_RECVRVCINFO`
- `SCTP_RECVNXTINFO`
- `SCTP_DEFAULT_SNDINFO`
- `SCTP_GETASSOC_NUMBER`
- `SCTP_GET_ASSOC_ID_LIST`

A few of the `opt` parameters are described in detail below:

<code>SCTP_RTOINFO</code>	Returns the protocol parameters that are used to initialize and bind the retransmission timeout (RTO) tunable. The protocol parameters use the following structure:
---------------------------	---

```

struct sctp_rtoinfo {
    sctp_assoc_t    srto_assoc_id;
    uint32_t       srto_initial;
    uint32_t       srto_max;
    uint32_t       srto_min;
};

```

`srto_assoc_id` The calling program provides this value, which specifies the association of interest.

`srto_initial` This value is the initial RTO value.

`srto_max` This value is the maximum RTO value.

`srto_min` This value is the minimum RTO value.

SCTP_ASSOCINFO

Returns the association-specific parameters. The parameters use the following structure:

```

struct sctp_assocparams {
    sctp_assoc_t    sasoc_assoc_id;
    uint16_t       sasoc_asocmaxrxt;
    uint16_t       sasoc_number_peer_destinations;
    uint32_t       sasoc_peer_rwnd;
    uint32_t       sasoc_local_rwnd;
    uint32_t       sasoc_cookie_life;
};

```

`sasoc_assoc_id`
The calling program provides this value, which specifies the association of interest.

`sasoc_asocmaxrxt`
This value specifies the maximum retransmission count for the association.

`sasoc_number_peer_destinations`
This value specifies the number of addresses that the peer has.

`sasoc_peer_rwnd`
This value specifies the current value of the peer's receive window.

`sasoc_local_rwnd`
This value specifies the last reported receive window that the peer transmitted to.

`sasoc_cookie_life`
The value specifies the lifetime of the association's cookie. The value is used when issuing cookies.

All parameters that use time values are measured in milliseconds.

SCTP_DEFAULT_SEND_PARAM	Returns the default set of parameters that a call to the <code>sendto(3SOCKET)</code> function uses on this association. The parameters use the following structure: <pre>struct sctp_sndrcvinfo { uint16_t sinfo_stream; uint16_t sinfo_ssn; uint16_t sinfo_flags; uint32_t sinfo_ppid; uint32_t sinfo_context; uint32_t sinfo_timetolive; uint32_t sinfo_tsn; uint32_t sinfo_cumtsn; sctp_assoc_t sinfo_assoc_id; };</pre> <p><code>sinfo_stream</code> This value specifies the default stream for the <code>sendmsg()</code> call.</p> <p><code>sinfo_ssn</code> This value is always zero.</p>
<code>sinfo_flags</code>	This value contains the default flags for the <code>sendmsg()</code> call. This flag can take on the following values: <ul style="list-style-type: none">▪ MSG_UNORDERED▪ MSG_ADDR_OVER▪ MSG_ABORT▪ MSG_EOF▪ MSG_PR_SCTP
<code>sinfo_ppid</code>	This value is the default payload protocol identifier for the <code>sendmsg()</code> call.
<code>sinfo_context</code>	This value is the default context for the <code>sendmsg()</code> call.
<code>sinfo_timetolive</code>	This value specifies a time period in milliseconds. After this time period has passed, the message expires if its transmission has not begun. A value of zero indicates that the message does not expire. If the MSG_PR_SCTP flag is set, the message expires when its transmission has not successfully completed within the time period specified in <code>sinfo_timetolive</code> .
<code>sinfo_tsn</code>	This value is always zero.
<code>sinfo_cumtsn</code>	This value is always zero.
<code>sinfo_assoc_id</code>	This value is filled in by the calling application. This value specifies the association of interest.
SCTP_PEER_ADDR_PARAMS	Returns the parameters for a specified peer address. The parameters use the following structure:

```

struct sctp_paddrparams {
    sctp_assoc_t      spp_assoc_id;
    struct sockaddr_storage spp_address;
    uint32_t          spp_hbinterval;
    uint16_t          spp_pathmaxrxt;
    uint32_t          spp_pathmtu;
    uint32_t          spp_flags;
    uint32_t          spp_ipv6_flowlabel;
    uint8_t           spp_ipv4_tos;
};

```

spp_assoc_id

The calling program provides this value, which specifies the association of interest.

spp_address

This value specifies the peer's address of interest.

spp_hbinterval

This value specifies the heartbeat interval in milliseconds.

spp_pathmaxrxt

This value specifies the maximum number of retransmissions to attempt on an address before considering the address unreachable.

spp_pathmtu

The current path MTU of the peer address. It is the number of bytes available in an SCTP packet for chunks. Providing a value of 0 does not change the current setting. If a positive value is provided and `SPP_PMTUD_DISABLE` is set in the `spp_flags`, the given value is used as the path MTU. If `SPP_PMTUD_ENABLE` is set in the `spp_flags`, then the `spp_pathmtu` field is ignored.

spp_ipv6_flowlabel

This field is used in conjunction with the `SPP_IPV6_FLOWLABEL` flag. This setting has precedence over any other IPv6 layer setting.

spp_flags

The `spp_flags` flags are used to control various features on an association. The flag field is a bit mask which may contain one of the following options:

- `SPP_HB_ENABLE` – Enable heartbeats on the specified address.

- **SPP_HB_DISABLE** – Disable heartbeats on the specified address. **SPP_HB_ENABLE** and **SPP_HB_DISABLE** are mutually exclusive, only one of the two should be specified. Enabling both fields will result in undetermined results.
- **SPP_HB_DEMAND** – Request a user initiated heartbeat to be made immediately. This option should not be used in conjunction with a wildcard address.
- **SPP_HB_TIME_IS_ZERO** – Specifies that the time for heartbeat delay is to be set to the value of 0 milliseconds.
- **SPP_PMTUD_ENABLE** – Enable PMTU discovery on the specified address.
- **SPP_PMTUD_DISABLE** – Disable PMTU discovery on the specified address. If the address field is empty then all addresses on the association are affected. **SPP_PMTUD_ENABLE** and **SPP_PMTUD_DISABLE** options are mutually exclusive. Enabling both options will result in undetermined results.
- **SPP_IPV6_FLOWLABEL** – Enables the setting of the IPV6 flowlabel value. The value is obtained from the `spp_ipv6_flowlabel` field. Upon retrieval, this flag will be set to indicate that the `ipv6_flowlabel` field has a valid value returned. If a specific destination address is set in the `spp_address` field, then the value of the address is returned. If only an association is specified and no address is specified, then the association's default flowlabel is returned. If neither an association nor a destination is specified, then the socket's default flowlabel is returned. For non-IPV6 sockets, this flag is left empty.
- **SPP_IPV4_TOS** – Setting this flag enables the setting of the IPV4 TOS value associated with either the association or a specific address. The value is obtained from the `spp_ipv4_tos` field. Upon retrieval, this flag will be set to indicate that the `spp_ipv4_tos` field has a valid value returned. If a specific destination address in the `spp_address` field is set when called, then the TOS value of the specific destination address returned. If only an association is specified then the default TOS of the association is returned. If neither an association nor a destination is specified, then the default TOS value of the socket is returned.

SCTP_STATUS

Returns the current status information about the association. The parameters use the following structure:

```
struct sctp_status {
    sctp_assoc_t      sstat_assoc_id;
    int32_t          sstat_state;
    uint32_t         sstat_rwnd;
    uint16_t         sstat_unackdata;
    uint16_t         sstat_penddata;
    uint16_t         sstat_instrms;
    uint16_t         sstat_outstrms;
    uint32_t         sstat_fragmentation_point;
    struct sctp_paddrinfo sstat_primary;
};
```

sstat_assoc_id

The calling program provides this value, which specifies the association of interest.

sstat_state

This value is the association's current state. The association can take on the following states:

SCTP_IDLE	The SCTP endpoint does not have any association associated with it. Immediately after the call to the <code>socket()</code> function opens an endpoint, or after the endpoint closes, the endpoint is in this state.
SCTP_BOUND	An SCTP endpoint is bound to one or more local addresses after calling the <code>bind()</code> .
SCTP_LISTEN	This endpoint is waiting for an association request from any remote SCTP endpoint.
SCTP_COOKIE_WAIT	This SCTP endpoint has sent an INIT chunk and is waiting for an INIT-ACK chunk.
SCTP_COOKIE_ECHOED	This SCTP endpoint has echoed the cookie that it

	received from its peer's INIT-ACK chunk back to the peer.
SCTP_ESTABLISHED	This SCTP endpoint can exchange data with its peer.
SCTP_SHUTDOWN_PENDING	This SCTP endpoint has received a SHUTDOWN primitive from its upper layer. This endpoint no longer accepts data from its upper layer.
SCTP_SHUTDOWN_SEND	An SCTP endpoint that was in the SCTP_SHUTDOWN_PENDING state has sent a SHUTDOWN chunk to its peer. The SHUTDOWN chunk is sent only after all outstanding data from this endpoint to its peer is acknowledged. When this endpoint's peer sends a SHUTDOWN ACK chunk, this endpoint sends a SHUTDOWN COMPLETE chunk and the association is considered closed.
SCTP_SHUTDOWN_RECEIVED	An SCTP endpoint has received a SHUTDOWN chunk from its peer. This endpoint no longer accepts new data from its user.
SCTP_SHUTDOWN_ACK_SEND	An SCTP endpoint in the SCTP_SHUTDOWN_RECEIVED state has sent the SHUTDOWN ACK chunk to its peer. The endpoint only sends the SHUTDOWN ACK chunk after the peer acknowledges all outstanding data from this endpoint. When this endpoint's peer sends a

SHUTDOWN COMPLETE chunk,
the association is closed.

`sstat_rwnd`

This value is the association peer's current receive window.

`sstat_unackdata`

This value is the number of unacknowledged DATA chunks.

`sstat_penddata`

This value is the number of DATA chunks that are awaiting receipt.

`sstat_instrms`

This value is the number of inbound streams.

`sstat_outstrms`

This value is the number of outbound streams.

`sstat_fragmentation_point`

If the combined size of the message, the SCTP headers, and the IP headers exceeds the value of `sstat_fragmentation_point`, the message fragments. This value is equal to the Path Maximum Transmission Unit (P-MTU) for the packet's destination address

`sstat_primary`

This value contains information about the primary peer address. This information uses the following structure:

```
struct sctp_paddrinfo {
    sctp_assoc_t          spinfo_assoc_id;
    struct sockaddr_storage spinfo_address;
    int32_t               spinfo_state;
    uint32_t              spinfo_cwnd;
    uint32_t              spinfo_srtt;
    uint32_t              spinfo_rto;
    uint32_t              spinfo_mtu;
};
```

`spinfo_assoc_id` The calling program provides this value, which specifies the association of interest.

`spinfo_address` This value is the primary peer's address.

`spinfo_state` This value can take on either of the two values `SCTP_ACTIVE` or `SCTP_INACTIVE`.

<code>spinfo_cwnd</code>	This value is the congestion window of the peer address.
<code>spinfo_srtt</code>	This value is the current smoothed round-trip time calculation for the peer address. This value is expressed in milliseconds.
<code>spinfo_rto</code>	This value is the current retransmission timeout value for the peer address. This value is expressed in milliseconds.
<code>spinfo_mtu</code>	This value is the P-MTU for the peer address.

The `sctp_opt_info()` function returns 0 on success. The `sctp_opt_info()` function returns -1 on failure and sets the value of `errno` to the appropriate error code. If the file descriptor passed to the `sctp_opt_info()` in the `sock` parameter is invalid, the `sctp_opt_info()` function fails and returns `EBADF`. If the file descriptor passed to the `sctp_opt_info()` function in the `sock` parameter does not describe a socket, the `sctp_opt_info()` function fails and returns `ENOTSOCK`. If the association ID is invalid for a one-to-many style SCTP socket, the `sctp_opt_info()` function fails and sets the value of `errno` to `EINVAL`. If the input buffer length is too short for the option specified, the `sctp_opt_info()` function fails and sets the value of `errno` to `EINVAL`. If the address family for the peer's address is not `AF_INET` or `AF_INET6`, the `sctp_opt_info()` function fails and sets the value of `errno` to `EAFNOSUPPORT`.

sctp_rcvmsg()

```
ssize_t sctp_rcvmsg(int s, void *msg, size_t len, struct sockaddr *from,  
socklen_t *fromlen, struct sctp_sndrcvinfo *sinfo, int *msg_flags);
```

The `sctp_rcvmsg()` function enables receipt of a message from the SCTP endpoint specified by the `s` parameter. The calling program can specify the following attributes:

msg

This parameter is the address of the message buffer.

len

This parameter is the length of the message buffer.

from

This parameter is a pointer to an address that contains the sender's address.

fromlen

This parameter is the size of the buffer associated with the address in the *from* parameter.

sinfo

This parameter is only active if the calling program enables `sctp_data_io_events`. To enable `sctp_data_io_events`, call the `setsockopt()` function with the socket option `SCTP_EVENTS`. When `sctp_data_io_events` is enabled, the application receives the contents of the `sctp_sndrcvinfo` structure for each incoming message. This parameter is a pointer to a `sctp_sndrcvinfo` structure. The structure is populated upon receipt of the message.

msg_flags

This parameter contains any message flags that are present.

The `sctp_rcvmsg()` function returns the number of bytes it receives. The `sctp_rcvmsg()` function returns -1 when an error occurs.

If the file descriptor passed in the *s* parameter is not valid, the `sctp_rcvmsg()` function fails and sets the value of `errno` to `EBADF`. If the file descriptor passed in the *s* parameter does not describe a socket, the `sctp_rcvmsg()` function fails and sets the value of `errno` to `ENOTSOCK`. If the *msg_flags* parameter includes the value `MSG_OOB`, the `sctp_rcvmsg()` function fails and sets the value of `errno` to `EOPNOTSUPP`. If there is no established association, the `sctp_rcvmsg()` function fails and sets the value of `errno` to `ENOTCONN`.

sctp_sendmsg()

```
ssize_t sctp_sendmsg(int s, const void *msg, size_t len, const struct sockaddr
*to, socklen_t tolen, uint32_t ppid, uint32_t flags, uint16_t stream_no, uint32_t
timetolive, uint32_t context);
```

The `sctp_sendmsg()` function enables advanced SCTP features while sending a message from an SCTP endpoint.

<i>s</i>	This value specifies the SCTP endpoint that is sending the message.
<i>msg</i>	This value contains the message sent by the <code>sctp_sendmsg()</code> function.
<i>len</i>	This value is the length of the message. This value is expressed in bytes.
<i>to</i>	This value is the destination address of the message.
<i>tolen</i>	This value is the length of the destination address.
<i>ppid</i>	This value is the application-specified payload protocol identifier.
<i>stream_no</i>	This value is the target stream for this message.
<i>timetolive</i>	This value is the time period after which the message expires if it has not been successfully sent to the peer. This value is expressed in milliseconds.
<i>context</i>	This value is returned if an error occurs during the sending of the message.
<i>flags</i>	This value is formed from applying the logical operation OR in bitwise fashion on zero or more of the following flag bits:

MSG_UNORDERED

When this flag is set, the `sctp_sendmsg()` function delivers the message unordered.

MSG_ADDR_OVER

When this flag is set, the `sctp_sendmsg()` function uses the address in the `to` parameter instead of the association's primary destination address. This flag is only used with one-to-many style SCTP sockets.

MSG_ABORT

When this flag is set, the specified association aborts by sending an ABORT signal to its peer. This flag is only used with one-to-many style SCTP sockets.

MSG_EOF

When this flag is set, the specified association enters graceful shutdown. This flag is only used with one-to-many style SCTP sockets.

MSG_PR_SCTP

When this flag is set, the message expires when its transmission has not successfully completed within the time period specified in the `timetolive` parameter.

The `sctp_sendmsg()` function returns the number of bytes it sent. The `sctp_sendmsg()` function returns -1 when an error occurs.

If the file descriptor passed in the `s` parameter is not valid, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EBADF`. If the file descriptor passed in the `s` parameter does not describe a socket, the `sctp_sendmsg()` function fails and sets the value of `errno` to `ENOTSOCK`. If the `flags` parameter includes the value `MSG_OOB`, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EOPNOTSUPP`. If the `flags` parameter includes the values `MSG_ABORT` or `MSG_EOF` for a one-to-one style socket, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EOPNOTSUPP`. If there is no established association, the `sctp_sendmsg()` function fails and sets the value of `errno` to `ENOTCONN`. If the socket is shutting down, disallowing further writes, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EPIPE`. If the socket is nonblocking and the transmit queue is full, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EAGAIN`.

If the control message length is incorrect, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EINVAL`. If the specified destination address does not belong to the association, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EINVAL`. If the value of `stream_no` is outside the number of outbound streams that the association supports, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EINVAL`. If the address family of the specified destination address is not `AF_INET` or `AF_INET6`, the `sctp_sendmsg()` function fails and sets the value of `errno` to `EINVAL`.

sctp_send()

```
ssize_t sctp_send(int s, const void *msg, size_t len, const struct
sctp_sndrcvinfo *sinfo, int flags);
```

The `sctp_send()` function is usable by one-to-one and one-to-many style sockets. The `sctp_send()` function enables advanced SCTP features while sending a message from an SCTP endpoint.

s

This value specifies the socket created by the `socket()` function.

msg

This value contains the message sent by the `sctp_send()` function.

len

This value is the length of the message. This value is expressed in bytes.

sinfo

This value contains the parameters used to send the message. For a one-to-many style socket, this value can contain the association ID to which the message is being sent.

flags

This value is identical to the `flags` parameter in the `sendmsg()` function.

The `sctp_send()` function returns the number of bytes it sent. The `sctp_send()` function returns -1 when an error occurs.

If the file descriptor passed in the `s` parameter is not valid, the `sctp_send()` function fails and sets the value of `errno` to `EBADF`. If the file descriptor passed in the `s` parameter does not describe a socket, the `sctp_send()` function fails and sets the value of `errno` to `ENOTSOCK`. If the `sinfo_flags` field of the `sinfo` parameter includes the value `MSG_OOB`, the `sctp_send()` function fails and sets the value of `errno` to `EOPNOTSUPP`. If the `sinfo_flags` field of the `sinfo` parameter includes the values `MSG_ABORT` or `MSG_EOF` for a one-to-one style socket, the `sctp_send()` function fails and sets the value of `errno` to `EOPNOTSUPP`. If there is no established association, the `sctp_send()` function fails and sets the value of `errno` to `ENOTCONN`. If the socket is shutting down, disallowing further writes, the `sctp_send()` function fails and sets the value of `errno` to `EPIPE`. If the socket is nonblocking and the transmit queue is full, the `sctp_send()` function fails and sets the value of `errno` to `EAGAIN`.

If the control message length is incorrect, the `sctp_send()` function fails and sets the value of `errno` to `EINVAL`. If the specified destination address does not belong to the association, the `sctp_send()` function fails and sets the value of `errno` to `EINVAL`. If the value of `stream_no` is outside the number of outbound streams that the association supports, the `sctp_send()` function fails and sets the value of `errno` to `EINVAL`. If the address family of the specified destination address is not `AF_INET` or `AF_INET6`, the `sctp_send()` function fails and sets the value of `errno` to `EINVAL`.

sctp_sendv()

```
ssize_t sctp_sendv(int sd, const struct iovec *iov, int iovcnt, struct sockaddr
*addrs, int addrcnt, void *info, socklen_t infolen, unsigned int infotype, int flags);
```

The `sctp_sendv()` sends a message to an SCTP socket. The following attributes are specified:

<code>sd</code>	The socket descriptor.								
<code>iov</code>	The message to be sent. The data in the buffer are treated as one single user message.								
<code>iovcnt</code>	The number of elements in <code>iov</code> .								
<code>addrs</code>	An array of addresses to be used to set up an association or one single address to be used to send the message. Pass in NULL if the caller does not want to set up an association nor want to send the message to a specific address.								
<code>addrcnt</code>	The number of addresses in the <code>addrs</code> array.								
<code>info</code>	A pointer to the buffer containing the attribute associated with the message to be sent. The type is indicated by <code>info_type</code> parameter.								
<code>infolen</code>	The length in bytes of <code>info</code> .								
<code>infotype</code>	The type of the <code>info</code> buffer. The following values are defined: <table><tr><td><code>SCTP_SENDV_SNDINFO</code></td><td>The type of <code>info</code> is <code>struct sctp_sndinfo</code>.</td></tr><tr><td><code>SCTP_SENDV_PRINFO</code></td><td>The type of <code>info</code> is <code>struct sctp_prinfo</code>.</td></tr><tr><td><code>SCTP_SENDV_AUTHINFO</code></td><td>The type of <code>info</code> is <code>struct sctp_authinfo</code>. This type is not supported.</td></tr><tr><td><code>SCTP_SENDV_SPA</code></td><td>The type of <code>info</code> is <code>struct sctp_send_spa</code>.</td></tr></table>	<code>SCTP_SENDV_SNDINFO</code>	The type of <code>info</code> is <code>struct sctp_sndinfo</code> .	<code>SCTP_SENDV_PRINFO</code>	The type of <code>info</code> is <code>struct sctp_prinfo</code> .	<code>SCTP_SENDV_AUTHINFO</code>	The type of <code>info</code> is <code>struct sctp_authinfo</code> . This type is not supported.	<code>SCTP_SENDV_SPA</code>	The type of <code>info</code> is <code>struct sctp_send_spa</code> .
<code>SCTP_SENDV_SNDINFO</code>	The type of <code>info</code> is <code>struct sctp_sndinfo</code> .								
<code>SCTP_SENDV_PRINFO</code>	The type of <code>info</code> is <code>struct sctp_prinfo</code> .								
<code>SCTP_SENDV_AUTHINFO</code>	The type of <code>info</code> is <code>struct sctp_authinfo</code> . This type is not supported.								
<code>SCTP_SENDV_SPA</code>	The type of <code>info</code> is <code>struct sctp_send_spa</code> .								

The `sctp_sendv()` function provides an extensible way for an application to communicate different send attributes to the SCTP stack when sending a message. This function can also be used to set up an association. The `addrs` array is similar to the `addrs` array used by [“sctp_connectx\(\)” on page 175](#).

There are three types of attributes which can be used to describe a message to be sent. They are represented by `struct sctp_sndinfo`, `struct sctp_prinfo`, and `struct sctp_authinfo` which is currently not supported.

The following structure `sctp_sendv_spa` is defined to be used when more than one of the above attributes are needed to describe a message to be sent.

```
struct sctp_sendv_spa {
    uint32_t sendv_flags;
    struct sctp_sndinfo sendv_sndinfo;
    struct sctp_prinfo sendv_prinfo;
```

```

    struct sctp_authinfo sendv_authinfo;
};

```

The *sendv_flags* field holds a bitwise-OR of SCTP_SEND_SNDINFO_VALID, Sctp_SEND_PRINFO_VALID, and Sctp_SEND_AUTHINFO_VALID values, indicating whether the *sendv_sndinfo*, *sendv_prinfo*, and *sendv_authinfo* fields contain valid information.

The *sctp_sndinfo* structure is defined as follows:

```

struct sctp_sndinfo {
    uint16_t    snd_sid;
    uint16_t    snd_flags;
    uint32_t    snd_ppid;
    uint32_t    snd_context;
    sctp_assoc_t  snd_assoc_id;
};

```

where:

<i>snd_sid</i>	This value holds the stream number to send the message to. If a sender specifies an invalid stream number, an error value is returned and the call fails.
<i>snd_flags</i>	This field is a bit wise OR of the following flags: <ul style="list-style-type: none"> SCTP_UNORDERED his flag requests the unordered delivery of the message. SCTP_ADDR_OVER This flag requests the SCTP stack to override the primary destination address and send the message to the given address in <i>addrs</i>. Only one address can be given in this case. If this flag is not specified and <i>addrs</i> is not NULL, this call is treated as a connect request. This flag is applicable to one-to-many style sockets only. SCTP_ABORT Setting this flag causes the specified association to be aborted by sending an ABORT message to the peer. The ABORT message will contain an error cause User Initiated Abort with cause code 12. The specific information the cause of this error is provided in <i>msg_iov</i>. SCTP_EOF Setting this flag invokes the SCTP graceful shutdown procedures on the specified association. Graceful shutdown assures that all data queued by both endpoints is successfully transmitted before closing the association. SCTP_SENALL This flag requests that the message is sent to all associations that are currently established on the socket. This flag is applicable to one-to-many style sockets only.

- snd_ppid* An unsigned integer that is passed to the remote end in each user message (SCTP DATA chunk). The SCTP stack performs no byte order modification of this field. For example, if the DATA chunk has to contain a given value in network byte order, the SCTP user has to perform the `htonl(3SOCKET)` computation.
- snd_context* This value is an opaque 32 bit context datum. It is passed back to the caller if an error occurs on the transmission of the message and is retrieved with each undelivered message.
- snd_assoc_id* When sending a message, this field holds the identifier for the association which the message is sent to. When this call is used to set up an association, the association identifier of the newly created association is returned in this field. This field is applicable to one-to-many style sockets only.

The *sctp_prinfo* structure is defined as follows:

```
struct sctp_prinfo {
    uint16_t pr_policy;
    uint32_t pr_value;
};
```

where:

- pr_policy* This field specifies the partial reliability (PR-SCTP) policy that is used to send the message. If it is `SCTP_PR_SCTP_NONE`, the message is sent reliably (the default is normal send). If it is `SCTP_PR_SCTP_TTL`, timed reliability as defined in RFC 3758 is used. In this case, the lifetime is provided in *pr_value*.
- pr_value* The meaning of this field depends on the PR-SCTP policy specified by the *pr_policy* field. It is ignored when `SCTP_PR_SCTP_NONE` is specified. In case of `SCTP_PR_SCTP_TTL`, this field specifies the lifetime in milliseconds of the message.

When new *send* attributes are needed, new structures can be defined. Those new structures do not need to be based on any of the above defined structures.

The `struct sctp_sndinfo` attribute for one-to-many style sockets must always be used in order to specify the association the message is to be sent to. The only case where it is not needed is when this call is used to set up a new association.

The caller provides a list of addresses in the *addrs* parameter to set up an association. This function will behave like calling `sctp_connect()`, first using the list of addresses, and then calling `sendmsg()` with the given message and attributes. For an one-to-many style socket, if a `struct sctp_sndinfo` attribute is provided, the *snd_assoc_id* field must be 0. When this function returns, the *snd_assoc_id* field will contain the association identifier of the newly established association. The `struct sctp_sndinfo` attribute is not required to set up an association for

one-to-many style sockets. If this attribute is not provided, the caller can enable the `SCTP_ASSOC_CHANGE` notification and use the `SCTP_COMM_UP` message to find out the association identifier.

If the caller wants to send the message to a specific peer address (overriding the primary address), the caller can provide the specific address in the *addrs* parameter and provide a `struct sctp_sndinfo` attribute with the *snd_flags* field set to `SCTP_ADDR_OVER`.

This function can also be used to terminate an association. The caller provides an *sctp_sndinfo* attribute with the *snd_flags* set to `SCTP_EOF`. In this case, the length of the message would be zero. Sending a message using `sctp_sendv()` is atomic unless explicit EOR marking is enabled on the socket specified by *sd*.

Upon successful completion, the number of bytes sent is returned. Otherwise, -1 is returned and *errno* is set to indicate the error.

The following error values are defined:

<code>EADDRINUSE</code>	The address is already in use.
<code>EADDRNOTAVAIL</code>	No local address is available for this operation.
<code>EAFNOSUPPORT</code>	Addresses in the specified address family cannot be used with this socket.
<code>EBADF</code>	The <i>sd</i> parameter is not a valid file descriptor.
<code>ECONNREFUSED</code>	The attempt to connect was forcefully rejected. The calling program should close the socket descriptor using <code>close(2)</code> and issue another <code>socket(3)</code> call to obtain a new descriptor before making another attempt.
<code>EFAULT</code>	A parameter can not be accessed.
<code>EINTR</code>	The operation was interrupted by delivery of a signal before any data could be buffered to be sent.
<code>EINVAL</code>	A parameter provided is invalid for this operation.
<code>EMSGSIZE</code>	The message is too large to be sent all at once.
<code>ENETUNREACH</code>	The network is not reachable from this host.
<code>ENOBUFS</code>	Insufficient memory is available to complete the operation.
<code>EOPNOTSUPP</code>	Operation not supported in this type of socket.
<code>EPIPE</code>	The peer end point has shutdown the association.
<code>ETIMEDOUT</code>	Attempt timed out.
<code>EWOULDBLOCK</code>	The socket is marked as non-blocking, and the requested operation would block.

sctp_recvv()

```
ssize_t sctp_recvv(int sd, const struct iovec *iov, int iovlen, struct sockaddr
*from, int fromlen, void *info, socklen_t infolen, unsigned int infotype, int flags);
```

The `sctp_recvv()` function provides an extensible way for the SCTP stack to pass up different SCTP attributes associated with a received message to an application. The following attributes are specified:

<i>sd</i>	The socket descriptor.								
<i>iov</i>	The scatter buffer containing the received message.								
<i>iovlen</i>	The number of elements in <i>iov</i> .								
<i>from</i>	A pointer to a buffer to be filled with the sender address of the received message.								
<i>fromlen</i>	The size of the <i>from</i> buffer. Upon return, it is set to the actual size of the sender's address.								
<i>info</i>	A pointer to the buffer containing the attributes of the received message. The type of structure is indicated by <i>infotype</i> parameter.								
<i>infolen</i>	The length in bytes of <i>info</i> buffer. Upon return, it is set to the actual size of the returned <i>info</i> buffer.								
<i>infotype</i>	The type of the info buffer. The following values are defined: <table> <tr> <td>SCTP_RECVM_NOINFO</td> <td>If both SCTP_RECVRVINFO and SCTP_RECVMXTINFO options are not enabled, no attribute will be returned. If only the SCTP_RECVMXTINFO option is enabled but there is no next message in the buffer, there will also no attribute be returned. In these cases, <i>infotype</i> will be set to SCTP_RECVM_NOINFO.</td> </tr> <tr> <td>SCTP_RECVM_RVINFO</td> <td>The type of <i>info</i> is <code>struct sctp_rcvinfo</code> and the attribute is about the received message.</td> </tr> <tr> <td>SCTP_RECVM_NXTINFO</td> <td>The type of <i>info</i> is <code>struct sctp_nxtinfo</code> and the attribute is about the next message in receive buffer. This is the case when only the SCTP_RECVMXTINFO option is enabled and there is a next message in the buffer.</td> </tr> <tr> <td>SCTP_RECVM_RN</td> <td>The type of <i>info</i> is <code>struct sctp_recvm_rn</code>. The <i>recvm_rcvinfo</i> field is the attribute of the received message and the <i>recvm_nxtinfo</i> field is the attribute of the next message in buffer. This is the case when both SCTP_RECVRVINFO and SCTP_RECVMXTINFO options are enabled and there is a next message in the receive buffer.</td> </tr> </table>	SCTP_RECVM_NOINFO	If both SCTP_RECVRVINFO and SCTP_RECVMXTINFO options are not enabled, no attribute will be returned. If only the SCTP_RECVMXTINFO option is enabled but there is no next message in the buffer, there will also no attribute be returned. In these cases, <i>infotype</i> will be set to SCTP_RECVM_NOINFO.	SCTP_RECVM_RVINFO	The type of <i>info</i> is <code>struct sctp_rcvinfo</code> and the attribute is about the received message.	SCTP_RECVM_NXTINFO	The type of <i>info</i> is <code>struct sctp_nxtinfo</code> and the attribute is about the next message in receive buffer. This is the case when only the SCTP_RECVMXTINFO option is enabled and there is a next message in the buffer.	SCTP_RECVM_RN	The type of <i>info</i> is <code>struct sctp_recvm_rn</code> . The <i>recvm_rcvinfo</i> field is the attribute of the received message and the <i>recvm_nxtinfo</i> field is the attribute of the next message in buffer. This is the case when both SCTP_RECVRVINFO and SCTP_RECVMXTINFO options are enabled and there is a next message in the receive buffer.
SCTP_RECVM_NOINFO	If both SCTP_RECVRVINFO and SCTP_RECVMXTINFO options are not enabled, no attribute will be returned. If only the SCTP_RECVMXTINFO option is enabled but there is no next message in the buffer, there will also no attribute be returned. In these cases, <i>infotype</i> will be set to SCTP_RECVM_NOINFO.								
SCTP_RECVM_RVINFO	The type of <i>info</i> is <code>struct sctp_rcvinfo</code> and the attribute is about the received message.								
SCTP_RECVM_NXTINFO	The type of <i>info</i> is <code>struct sctp_nxtinfo</code> and the attribute is about the next message in receive buffer. This is the case when only the SCTP_RECVMXTINFO option is enabled and there is a next message in the buffer.								
SCTP_RECVM_RN	The type of <i>info</i> is <code>struct sctp_recvm_rn</code> . The <i>recvm_rcvinfo</i> field is the attribute of the received message and the <i>recvm_nxtinfo</i> field is the attribute of the next message in buffer. This is the case when both SCTP_RECVRVINFO and SCTP_RECVMXTINFO options are enabled and there is a next message in the receive buffer.								

flags Flag for receive as in `recvmsg(3SOCKET)`. On return, its value will be different from what was set in to the call. It has the same value as `rcv_flags`.

There are two types of attributes which can be returned by the call to `sctp_rcvv()`:

- The attribute of the received message and the attribute of the next message in *receive* buffer. The caller enables the `SCTP_RECVRCVINFO` and `SCTP_RECVNXTINFO` socket option to receive these attributes respectively.

Attributes of the received message are returned in `struct sctp_rcvinfo` and attributes of the next message are returned in the structure `sctp_nxtinfo`. If both options are enabled, both attributes are returned using the following structure.

```
struct sctp_rcvv_rn {
    struct sctp_rcvinfo rcvv_rcvinfo;
    struct sctp_nxtinfo rcvv_nxtinfo;
};
```

The `sctp_rcvinfo` structure is defined as follows:

```
struct sctp_rcvinfo {
    uint16_t rcv_sid;
    uint16_t rcv_ssn;
    uint16_t rcv_flags;
    uint32_t rcv_ppid;
    uint32_t rcv_tsn;
    uint32_t rcv_cumtsn;
    uint32_t rcv_context;
    sctp_assoc_t rcv_assoc_id;
};
```

where:

<code>rcv_info</code>	The stream number of the received message.
<code>rcv_ssn</code>	The stream sequence number that the peer endpoint assigned to the DATA chunk of this message. For fragmented messages, this is the same number for all deliveries of the message (if more than one <code>sctp_rcvv()</code> is needed to read the message).
<code>rcv_flags</code>	May be set to <code>SCTP_UNORDERED</code> when the message was sent unordered.
<code>rcv_ppid</code>	This value is the same information that is passed by the peer socket to its SCTP stack. The SCTP stack performs no byte order modification of this field.
<code>rcv_tsn</code>	The transmission sequence number that the peer endpoint assigned to the received message.
<code>rcv_cumtsn</code>	The current cumulative transmission sequence number of the association known to the SCTP stack.

<code>rcv_assoc_id</code>	The association identifier of the association of the received message. This field applies only to a one-to-many style socket.
<code>rcv_context</code>	This value is an opaque 32 bit context datum that was set by the caller with the <code>SCTP_CONTEXT</code> socket option. This value is passed back to the upper layer if an error occurs on the transmission of a message and is retrieved with each undelivered message.

The `sctp_nxtinfo` structure is defined as follows:

```
struct sctp_nxtinfo {
    uint16_t nxt_sid;
    uint16_t nxt_flags;
    uint32_t nxt_ppid;
    size_t   nxt_length;
    sctp_assoc_t nxt_assoc_id;
};
```

where:

<i>nxt_sid</i>	The stream number of the next message.						
<i>flags</i>	This field can contain any of the following flags and is composed of a bitwise-OR of the following values: <table><tr><td><code>SCTP_UNORDERED</code></td><td>The next message was sent unordered.</td></tr><tr><td><code>SCTP_COMPLETE</code></td><td>The entire message has been received and is in the socket buffer. This flag has special implications with respect to the <i>nxt_length</i> field.</td></tr><tr><td><code>SCTP_NOTIFICATION</code></td><td>The next message is not a user message but instead is a notification.</td></tr></table>	<code>SCTP_UNORDERED</code>	The next message was sent unordered.	<code>SCTP_COMPLETE</code>	The entire message has been received and is in the socket buffer. This flag has special implications with respect to the <i>nxt_length</i> field.	<code>SCTP_NOTIFICATION</code>	The next message is not a user message but instead is a notification.
<code>SCTP_UNORDERED</code>	The next message was sent unordered.						
<code>SCTP_COMPLETE</code>	The entire message has been received and is in the socket buffer. This flag has special implications with respect to the <i>nxt_length</i> field.						
<code>SCTP_NOTIFICATION</code>	The next message is not a user message but instead is a notification.						
<i>ppid</i>	This value is the same information that was passed by the peer socket to its SCTP stack when sending the next message. The SCTP stack performs no byte order modification of this field.						
<i>length</i>	The length of the message currently received in the socket buffer. This might not be the entire length of the next message since a partial delivery may be in progress. This field represents the entire next message size only if the flag <code>SCTP_COMPLETE</code> is set in the <i>nxt_flags</i> field.						
<i>assoc_id</i>	The association identifier of the association of the next message. This field applies only to a one-to-many style socket.						

The following error values are defined for `sctp_recv()`:

<code>EBADF</code>	The <i>sd</i> parameter is not a valid file descriptor.
<code>EFAULT</code>	A parameter can not be accessed.

EINTR	The operation was interrupted by delivery of a signal before any data could be buffered to be sent or the operation was interrupted by delivery of a signal before any data is available to be received.
EINVAL	A parameter provided is invalid for this operation.
ENOBUFS	Insufficient memory is available to complete the operation.
EWOULDBLOCK	The socket is marked as non-blocking and the requested operation would get blocked.

sctp_connectx()

```
int sctp_connectx(int sd, struct sockaddr *addrs, int addrcnt, sctp_assoc_t *aid);
```

The `sctp_connectx()` requests an SCTP association to be made on a socket. This is similar to `connect(3SOCKET)` except that an array of peer addresses can be given.

Much like `sctp_bindx(3SOCKET)`, this function allows a caller to specify multiple addresses at which a peer can be reached. The SCTP stack tries each addresses in the array in a round robin fashion to set up the association. Note that the list of addresses passed in is only used for setting up the association. It does not necessarily equal the set of addresses the peer uses for the resulting association. If the caller wants to find out the set of peer addresses, the caller must use `sctp_getpaddrs(3SOCKET)` to retrieve them after the association has been set up.

The following attributes are specified:

<code>sd</code>	The socket descriptor.
<code>addrs</code>	If <code>sd</code> is an IPv4 socket, <code>addrs</code> should be an array of <code>sockaddr_in</code> structures containing IPv4 addresses. If <code>sd</code> is an IPv6 socket, <code>addrs</code> should be an array of <code>sockaddr_in6</code> structures containing IPv6 or IPv4-mapped IPv6 addresses.
<code>addrcnt</code>	The number of addresses in the array <code>addrs</code> .
<code>aid</code>	If the call to <code>sctp_connectx()</code> function returns successfully, the association identifier for the newly created association is returned in <code>aid</code> . This parameter is applicable only to one-to-many style SCTP sockets.

The following error values are defined for `sctp_connectx()`:

EADDRINUSE	The address is already in use.
EADDRNOTAVAIL	No local address is available for this operation.
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	The <code>sd</code> parameter is not a valid file descriptor.

ECONNREFUSED	The attempt to connect was forcefully rejected. The calling program should use <code>connect(3SOCKET)</code> to close the socket descriptor, and issue another <code>socket(3SOCKET)</code> call to obtain a new descriptor before making another attempt.
EFAULT	A parameter can not be accessed.
EINTR	The connect attempt was interrupted before it is completed. The attempt will be established asynchronously.
EINVAL	A parameter provided is invalid for this operation.
ENOBUFS	Insufficient memory is available to complete the operation.
EWOULDBLOCK	The socket is marked as non-blocking and the requested operation would get blocked.
ETIMEDOUT	The attempt timed out.
EOPNOTSUPP	The operation is not supported in this type of socket.

sctp_getladdrs()

The `sctp_getladdrs()` function returns all locally bound addresses on a socket. The syntax for the `sctp_getladdrs()` function is as follows:

```
int sctp_getladdrs(int sock, sctp_assoc_t id, void **addrs);
```

When the `sctp_getladdrs()` function returns successfully, the value of `addrs` points to a dynamically allocated packed array of `sockaddr` structures. The `sockaddr` structures are of the appropriate type for each local address. The calling application uses the `sctp_freeladdrs()` function to free the memory. The value of the `addrs` parameter must not be NULL.

If the socket referenced by the `sd` parameter is an IPv4 socket, the `sctp_getladdrs()` function returns IPv4 addresses. If the socket referenced by the `sd` parameter is an IPv6 socket, the `sctp_getladdrs()` function returns a mix of IPv4 or IPv6 addresses as appropriate.

When the `sctp_getladdrs()` function is invoked for a one-to-many style socket, the value of the `id` parameter specifies the association to query. The `sctp_getladdrs()` function ignores the `id` parameter when the function is operating on a one-to-one socket.

When the value of the `id` parameter is zero, the `sctp_getladdrs()` function returns locally bound addresses without regard to any particular association. When the `sctp_getladdrs()` function returns successfully, it reports the number of local addresses bound to the socket. If the socket is unbound, the `sctp_getladdrs()` function returns 0 and the value of `*addrs` is undefined. If an error occurs, the `sctp_getladdrs()` function returns -1 and the value of `*addrs` is undefined.

sctp_freeladdrs()

The `sctp_freeladdrs()` function frees all of the resources that were allocated by a previous call to the `sctp_getladdrs()`. The syntax for the `sctp_freeladdrs()` function is as follows:

```
void sctp_freeladdrs(void *addrs);
```

The `*addrs` parameter is an array that contains the peer addresses that are returned by the `sctp_getladdrs()` function.

sctp_getpaddrs()

The `sctp_getpaddrs()` function returns all peer addresses in an association.

```
int sctp_getpaddrs(int sock, sctp_assoc_t id, void **addrs);
```

When the `sctp_getpaddrs()` function returns successfully, the value of the `**addrs` parameter points to a dynamically allocated packed array of `sockaddr` structures of the appropriate type for each address. The calling thread frees the memory with the `sctp_freepaddrs()` function. The `**addrs` parameter cannot have a value of `NULL`. If the socket descriptor given in `sock` is for an IPv4 socket, the `sctp_getpaddrs()` function returns IPv4 addresses. If the socket descriptor given in `sock` is for an IPv6 socket, the `sctp_getpaddrs()` function returns a mix of IPv4 and IPv6 addresses. For one-to-many style sockets, the `id` parameter specifies the association to query. The `sctp_getpaddrs()` function ignores the `id` parameter for one-to-one style sockets. When the `sctp_getpaddrs()` function returns successfully, it returns the number of peer addresses in the association. If there is no association on this socket, the `sctp_getpaddrs()` function returns 0 and the value of the `**addrs` parameter is undefined. If an error occurs, the `sctp_getpaddrs()` function returns -1 and the value of the `**addrs` parameter is undefined.

If the file descriptor passed to the `sctp_getpaddrs()` function in the `sock` parameter is invalid, the `sctp_getpaddrs()` function fails and returns `EBADF`. If the file descriptor passed to the `sctp_getpaddrs()` function in the `sock` parameter does not describe a socket, the `sctp_getpaddrs()` function fails and returns `ENOTSOCK`. If the file descriptor passed to the `sctp_getpaddrs()` function in the `sock` parameter describes a socket that is not connected, the `sctp_getpaddrs()` function fails and returns `ENOTCONN`.

sctp_freepaddrs()

The `sctp_freepaddrs()` function frees all of the resources that were allocated by a previous call to the `sctp_getpaddrs()`. The syntax for the `sctp_freepaddrs()` function is as follows:

```
void sctp_freepaddrs(void *addrs);
```

The `*addrs` parameter is an array that contains the peer addresses that are returned by the `sctp_getpaddrs()` function.

Branched-off Association

Applications can branch an established association on a one-to-many style socket into a separate socket and file descriptor. A separate socket and file descriptor is useful for applications that have a number of sporadic message senders or receivers that need to remain under the original one-to-many style socket. The application branches off associations that carry high volume data traffic into separate socket descriptors. The application uses the `sctp_peekoff()` call to branch off an association into a separate socket. The new socket is a one-to-one style socket. The syntax for the `sctp_peekoff()` function is as follows:

```
int sctp_peekoff(int sock, sctp_assoc_t id);
```

sock

The original one-to-many style socket descriptor returned from the `socket()` system call.

id

The identifier of the association to branch off to a separate file descriptor.

The `sctp_peekoff()` function fails and returns `EOPNOTSUPP` if the socket descriptor passed in the *sock* parameter is not a one-to-many style SCTP socket. The `sctp_peekoff()` function fails and returns `EINVAL` if the value of *id* is zero or if the value of *id* is greater than the maximum number of associations for the socket descriptor passed in the *sock* parameter. The `sctp_peekoff()` function fails and returns `EMFILE` if the function fails to create a new user file descriptor or file structure.

Code Examples of SCTP Use

This section details two uses of SCTP sockets.

EXAMPLE 7-17 SCTP Echo Client

```
/*
 * Copyright (c) 2012, Oracle and/or its affiliates. All rights reserved.
 */

/*
 * IPv4 echo client.
 */

/* To enable socket features used for SCTP socket. */
#define _XPG4_2
#define __EXTENSIONS__

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

EXAMPLE 7-17 SCTP Echo Client (Continued)

```

#include <netinet/sctp.h>
#include <errno.h>
#define BUFLLEN 2048
#define SERVER_PORT 5000
#define MAX_STREAM 64

static void
usagBe(char *a0)
{
    fprintf(stderr, "Usage: %s <server address>\n", a0);
}
static void
print_notif(char *buf)
{
    union sctp_notification *snp;
    struct sctp_assoc_change *sac;
    snp = (union sctp_notification *)buf;
    /* We only subscribe the association change event. */
    if (snp->sn_header.sn_type != SCTP_ASSOC_CHANGE) {
        fprintf(stderr, "unexpected notification type: %d\n",
            snp->sn_header.sn_type);
        exit(1);
    }
    sac = &snp->sn_assoc_change;
    printf("[ Receive association change event: state = %hu," error = %hu,"
    instr = %hu, outstr = %hu ]\n", sac->sac_state,
        sac->sac_error, sac->sac_inbound_streams,
        sac->sac_outbound_streams);
}

/*
 * Read from the network.
 */
static void
readit(void *vfdp)
{
    int fd;
    ssize_t n;
    char buf[BUFLLEN];
    struct iovec iov[1];
    int flags;
    socklen_t info_len;
    uint_t info_type;
    struct sctp_rcvinfo info;
    union sctp_notification *snp;

    pthread_setcanceltype(PTHREAD_CANCEL_ASYNCHRONOUS, NULL);
    fd = *(int *)vfdp;

    /* Initialize the iov for receiving */
    memset(buf, 0, BUFLLEN);
    iov->iov_base = buf;
    iov->iov_len = BUFLLEN;

    info_len = sizeof (info);
    info_type = 0;

```

EXAMPLE 7-17 SCTP Echo Client (Continued)

```

flags = 0;
while ((n = sctp_rcvv(fd, iov, 1, NULL, NULL, &info,
&info_len, &info_type, &flags)) > 0) {
    /* Intercept notifications here */
    if (flags & MSG_NOTIFICATION) {
        print_notif(buf);
        continue;
    }
    /* The message should be accompanied by sctp_rcvinfo. */
    if (info_type != SCTP_RECVV_RCVINFO) {
        fprintf(stderr, "no sctp_rcvinfo attached\n");
        exit(1);
    }
    printf("[ Receive echo (%u bytes): stream = %hu, ssn = %hu," "tsn = %hu,
flags = %hx, ppid = %u ]\n", n, info.rcv_sid, info.rcv_ssn, info.rcv_tsn,
info.rcv_flags, info.rcv_ppid);
    flags = 0;
    info_len = sizeof (info);
}
if (n < 0) {
    perror("sctp_rcvv");
    exit(1);
}
close(fd);
exit(0);
}

```

```

static void
echo(struct sockaddr_in *addr, int addrcnt)
{
    int          fd;
    uchar_t      buf[BUFLLEN];
    ssize_t      n;
    int          perr;
    pthread_t    tid;
    struct iovec  iov[1];
    int          ret, on;
    struct sctp_sndinfo  info;
    struct sctp_initmsg  initmsg;
    struct sctp_event    event;

    /* Create a one-one SCTP socket */
    if ((fd = socket(AF_INET, SOCK_STREAM, IPPROTO_SCTP)) == -1) {
        perror("socket");
        exit(1);
    }
    /*
     * We are interested in association change events and we want
     * to get sctp_rcvinfo in each receive.
     */
    event.se_assoc_id = 0; /* Ignored for one-one SCTP socket */
    event.se_type = SCTP_ASSOC_CHANGE;
    event.se_on = 1;
    ret = setsockopt(fd, IPPROTO_SCTP, SCTP_EVENT, &event, sizeof (event));
    if (ret < 0) {
        perror("setsockopt SCTP_EVENT");
    }
}

```

EXAMPLE 7-17 SCTP Echo Client (Continued)

```

    exit(1);
}
on = 1;
ret = setsockopt(fd, IPPROTO_SCTP, SCTP_RECVRCVINFO, &on, sizeof (on));
if (ret < 0) {
    perror("setsockopt SCTP_RECVRCVINFO");
    exit(1);
}

/*
 * Set the SCTP stream parameters to tell the other side when
 * setting up the association.
 */
memset(&initmsg, 0, sizeof (struct sctp_initmsg));
initmsg.sinit_num_ostreams = MAX_STREAM;
initmsg.sinit_max_instreams = MAX_STREAM;
initmsg.sinit_max_attempts = MAX_STREAM;
ret = setsockopt(fd, IPPROTO_SCTP, SCTP_INITMSG, &initmsg, sizeof (struct sctp_initmsg));
if (ret < 0) {
    perror("setsockopt SCTP_INITMSG");
    exit(1);
}

/* Now connect to the peer. */
if (sctp_connectx(fd, (struct sockaddr *)addrs, addrcnt, NULL) == -1) {
    perror("sctp_connectx");
    exit(1);
}

/* Initialize the struct sctp_sndinfo for sending. */
memset(&info, 0, sizeof (info));
/* Start sending to stream 0. */
info.snd_sid = 0;
/*
 * Note that the server is expected to echo back the snd_ppid value.
 * So we don't need to do any conversion here. But if the server needs
 * to understand this value, we need to do a htonl() on it so that the
 * server side can do a ntohl() to convert it back to the host byte
 * order.
 */
info.snd_ppid = 0;

/* Create a thread to receive network traffic. */
perr = pthread_create(&tid, NULL, (void (*)(void *))readit, &fd);
if (perr != 0) {
    fprintf(stderr, "pthread_create: %d\n", perr);
    exit(1);
}

iov->iov_base = buf;
/* Read from stdin and then send to the echo server. */
while ((n = read(fileno(stdin), buf, BUFLen)) > 0) {
    iov->iov_len = n;
    if (sctp_sendv(fd, iov, 1, NULL, 0, &info, sizeof (info),
        SCTP_SENDV_SNDINFO, 0) < 0) {
        perror("sctp_sendv");
    }
}

```

EXAMPLE 7-17 SCTP Echo Client (Continued)

```

        exit(1);
    }
    /* Send the next message to a different stream. */
    info.snd_sid = (info.snd_sid + 1) % MAX_STREAM;
    info.snd_ppid++;
}
pthread_cancel(tid);
close(fd);
}

static struct sockaddr_in *
setup_addrs(const char *name, int *addrcnt)
{
    int          num_addrs, i;
    int          error;
    struct hostent *hp;
    struct sockaddr_in *addrs;

    hp = getipnodebyname(name, AF_INET, AI_DEFAULT, &error);
    if (hp == NULL) {
        fprintf(stderr, "host %s not found\n", name);
        return (NULL);
    }
    for (num_addrs = 0; hp->h_addr_list[num_addrs] != NULL; num_addrs++)
        ;
    addrs = malloc((num_addrs) * sizeof (*addrs));
    if (addrs == NULL) {
        fprintf(stderr, "cannot allocate address list\n");
        return (NULL);
    }
    for (i = 0; i < num_addrs; i++) {
        addrs[i].sin_family = AF_INET;
        addrs[i].sin_addr.s_addr = *(ipaddr_t *)hp->h_addr_list[i];
        addrs[i].sin_port = htons(SERVER_PORT);
    }
    *addrcnt = num_addrs;
    return (addrs);
}

int
main(int argc, char **argv)
{
    struct sockaddr_in *addrs;
    int addrcnt;

    if (argc < 2) {
        usage(*argv);
        exit(1);
    }

    /* Find the host to connect to. */
    if ((addrs = setup_addrs(argv[1], &addrcnt)) == NULL)
        exit(1);
    echo(addrs, addrcnt);
    return (0);
}

```

EXAMPLE 7-18 SCTP Echo Server

```

/*
 * Copyright (c) 2012, Oracle and/or its affiliates. All rights reserved.
 */
/*
 * IPv4 echo server
 */
/* To enable socket features used for SCTP socket. */
#define _XPG4_2
#define __EXTENSIONS__

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

#define BUFLen 1024
#define SERVER_PORT 5000
#define MAX_STREAM 64
/*
 * Given an event notification, print out what it is.
 */
static void
handle_event(void *buf)
{
    struct sctp_assoc_change *sac;
    struct sctp_send_failed_event *ssfe;
    struct sctp_paddr_change *spc;
    struct sctp_remote_error *sre;
    union sctp_notification *snp;
    char addrbuf[INET6_ADDRSTRLEN];
    const char *ap;
    struct sockaddr_in *sin;

    snp = buf;
    switch (snp->sn_header.sn_type) {
    case SCTP_ASSOC_CHANGE:
        sac = &snp->sn_assoc_change;
        printf(">>> assoc_change: state=%hu, error=%hu, instr=%hu "
            "outstr=%hu\n", sac->sac_state, sac->sac_error,
            sac->sac_inbound_streams, sac->sac_outbound_streams);
        break;
    case SCTP_SEND_FAILED_EVENT:
        ssfe = &snp->sn_send_failed_event;
        printf(">>> sendfailed: len=%hu err=%d\n", ssfe->ssfe_length,
            ssfe->ssfe_error);
        break;
    case SCTP_PEER_ADDR_CHANGE:
        spc = &snp->sn_paddr_change;
        if (spc->spc_aaddr.ss_family != AF_INET) {
            fprintf(stderr, "getmsg: unexpected family %d\n", spc->spc_aaddr.ss_family);
            exit(1);
        }
    }
}

```

EXAMPLE 7-18 SCTP Echo Server (Continued)

```

    } else {
        sin = (struct sockaddr_in *)&spc->spc_addr;
        ap = inet_ntop(AF_INET, &sin->sin_addr, addrbuf, INET6_ADDRSTRLEN);
    }
    printf(">>> intf_change: %s state=%d, error=%d\n", ap,
        spc->spc_state, spc->spc_error);
    break;
case SCTP_REMOTE_ERROR:
    sre = &snp->sn_remote_error;
    printf(">>> remote_error: err=%hu len=%hu\n",
        ntohs(sre->sre_error), ntohs(sre->sre_length));
    break;
case SCTP_SHUTDOWN_EVENT:
    printf(">>> shutdown event\n");
    break;
default:
    printf(">>> unexpected type: %hu\n", snp->sn_header.sn_type);
    break;
}
}

/*
 * Receive a message from the network.
 */
static ssize_t
getmsg(int fd, struct iovec *iov, struct sctp_rcvinfo *info, int *flags)
{
    ssize_t      tot = 0, nr;
    size_t       buflen;
    socklen_t    info_len;
    uint_t       info_type;
    char         *buf;

    buf = iov->iov_base;
    buflen = iov->iov_len;

    /* Loop until a whole message is received. */
    for (;;) {
        info_len = sizeof (*info);
        memset(info, 0, sizeof (*info));
        *flags = 0;
        nr = sctp_rcvfv(fd, iov, 1, NULL, NULL, info, &info_len,
            &info_type, flags);
        if (nr <= 0) {
            /* EOF or error */
            iov->iov_base = buf;
            return (nr);
        }
        tot += nr;

        /* Whole message/notification is received, return it. */
        if (*flags & MSG_EOR || *flags & MSG_NOTIFICATION) {
            iov->iov_base = buf;
            /* Buffer may be realloc(). Return the new size. */
            iov->iov_len = buflen;
            return (tot);
        }
    }
}

```

EXAMPLE 7-18 SCTP Echo Server (Continued)

```

    }

    /* Only sctp_rcvinfo is expected. */
    if (info_type != SCTP_RECVV_RCVINFO) {
        fprintf(stderr, "unexpected info received: %d\n",
            info_type);
        iov->iov_base = buf;
        return (-1);
    }

    /* Maybe we need a bigger buffer, do realloc(). */
    if (buflen == tot) {
        buf = realloc(buf, buflen * 2);
        if (buf == NULL) {
            fprintf(stderr, "out of memory\n");
            exit(1);
        }
        buflen *= 2;
    }
    /* Set the next read offset */
    iov->iov_base = buf + tot;
    iov->iov_len = buflen - tot;
}

}

/*
 * The echo server.
 */
static void
echo(int fd)
{
    ssize_t      nr;
    size_t       buflen;
    int          flags;
    struct iovec  iov[1];
    struct sctp_rcvinfo  rinfo;
    struct sctp_sndinfo  sinfo;

    if ((iov->iov_base = malloc(BUFLEN)) == NULL) {
        fprintf(stderr, "out of memory\n");
        exit(1);
    }
    iov->iov_len = BUFLEN;

    memset(&sinfo, 0, sizeof (sinfo));

    /* Wait for something to echo */
    while ((nr = getmsg(fd, iov, &rinfo, &flags)) > 0) {
        /* Intercept notifications here */
        if (flags & SCTP_NOTIFICATION) {
            handle_event(iov->iov_base);
            continue;
        }

        printf(">>> got %u bytes on stream %hu: ", nr, rinfo.rcv_sid);
        fflush(stdout);
    }
}

```

EXAMPLE 7-18 SCTP Echo Server (Continued)

```

write(fileno(stdout), iov->iov_base, nr);
fflush(stdout);

/* The buffer may be realloc(), so get the new size. */
buflen = iov->iov_len;

/*
 * Echo the message back using the incoming info.
 *
 * Note that rcv_sid is in host byte order. But rcv_ppid is
 * what is stored by the peer. If both sides want to use this
 * value for communication (interpreting it on both sides),
 * the sender needs to do htonl() when setting snd_ppid. And
 * the receiver side needs to do ntohl() to convert rcv_ppid
 * back to the host byte order.
 */
sinfo.snd_sid = rinfo.rcv_sid;
sinfo.snd_ppid = rinfo.rcv_ppid;
iov->iov_len = nr;

if (sctp_sendv(fd, iov, 1, NULL, 0, &sinfo, sizeof (sinfo),
SCTP_SENDV_SNDINFO, 0) < 0) {
    fprintf(stderr, "sctp_sendv\n");
    exit(1);
}

/* Restore the original buffer size. */
iov->iov_len = buflen;
}
free(iov->iov_base);
close(fd);
}

static void
subscribe_event(int fd, uint16_t event)
{
    struct sctp_event    ev;
    int                  ret;

    ev.se_assoc_id = 0; /* Ignored for one-one SCTP socket */
    ev.se_type = event;
    ev.se_on = 1;
    ret = setsockopt(fd, IPPROTO_SCTP, SCTP_EVENT, &ev, sizeof (ev));
    if (ret < 0) {
        fprintf(stderr, "%s: setsockopt SCTP_EVENT: %d\n", strerror(errno), event);
        exit(1);
    }
}

/* List of events we are interested in. */
static uint16_t event_interested[] = {
    SCTP_ASSOC_CHANGE,
    SCTP_SEND_FAILED_EVENT,
    SCTP_PEER_ADDR_CHANGE,
    SCTP_REMOTE_ERROR,
    SCTP_SHUTDOWN_EVENT

```

EXAMPLE 7-18 SCTP Echo Server (Continued)

```

};

int main(void)
{
    int      lfd;
    int      cfd;
    int      onoff;
    int      i;
    struct sockaddr_in  sin[1];
    struct sctp_initmsg  initmsg;

    if ((lfd = socket(AF_INET, SOCK_STREAM, IPPROTO_SCTP)) == -1) {
        perror("socket");
        exit(1);
    }
    sin->sin_family = AF_INET;
    sin->sin_port = htons(SERVER_PORT);
    sin->sin_addr.s_addr = INADDR_ANY;
    if (bind(lfd, (struct sockaddr *)sin, sizeof (*sin)) == -1) {
        perror("bind");
        exit(1);
    }
    if (listen(lfd, 1) == -1) {
        perror("listen");
        exit(1);
    }

    (void) memset(&initmsg, 0, sizeof (struct sctp_initmsg));
    initmsg.sinit_num_ostreams = MAX_STREAM;
    initmsg.sinit_max_instreams = MAX_STREAM;
    initmsg.sinit_max_attempts = MAX_STREAM;
    if (setsockopt(lfd, IPPROTO_SCTP, SCTP_INITMSG, &initmsg,
        sizeof (struct sctp_initmsg)) < 0) {
        perror("SCTP_INITMSG");
        exit(1);
    }

    /* Subscribe to events. */
    for (i = 0; i < sizeof (event_interested) / sizeof (uint16_t); i++)
        subscribe_event(lfd, event_interested[i]);

    /* Wait for new associations */
    for (;;) {
        if ((cfd = accept(lfd, NULL, 0)) == -1) {
            perror("accept");
            exit(1);
        }
        /* Subscribe to interesting events for the new association. */
        for (i = 0; i < sizeof (event_interested) / sizeof (int); i++)
            subscribe_event(cfd, event_interested[i]);

        /* We want sctp_rcvinfo in each receive. */
        onoff = 1;
        i = setsockopt(cfd, IPPROTO_SCTP, SCTP_RECVRCVINFO, &onoff,
            sizeof (onoff));
        if (i < 0) {

```

EXAMPLE 7-18 SCTP Echo Server *(Continued)*

```
        perror("setsockopt SCTP_RECVRCVINFO");
        close(cfd);
        continue;
    }

    /* Echo back any and all data */
    echo(cfd);
}
}
```

Programming With XTI and TLI

This chapter describes the Transport Layer Interface (TLI) and the X/Open Transport Interface (XTI). Advanced topics such as asynchronous execution mode are discussed in [“Advanced XTI/TLI Topics” on page 193](#).

Some recent additions to XTI, such as scatter/gather data transfer, are discussed in [“Additions to the XTI Interface” on page 213](#).

The transport layer of the OSI model (layer 4) is the lowest layer of the model that provides applications and higher layers with end-to-end service. This layer hides the topology and characteristics of the underlying network from users. The transport layer also defines a set of services common to many contemporary protocol suites including the OSI protocols, Transmission Control Protocol and TCP/IP Internet Protocol Suite, Xerox Network Systems (XNS), and Systems Network Architecture (SNA).

TLI is modeled on the industry standard Transport Service Definition (ISO 8072). It also can be used to access both TCP and UDP. XTI and TLI are a set of interfaces that constitute a network programming interface. The Oracle Solaris operating system supports both interfaces, although XTI represents the future direction of this set of interfaces. The Oracle Solaris software implements XTI and TLI as a user library using the STREAMS I/O mechanism.

What Are XTI and TLI?

Note – The interfaces described in this chapter are multithread safe. This means that applications containing XTI/TLI interface calls can be used freely in a multithreaded application. Because these interface calls are not re-entrant, they do not provide linear scalability.



Caution – The XTI/TLI interface behavior has not been well specified in an asynchronous environment. Do not use these interfaces from signal handler routines.

TLI was introduced with AT&T System V, Release 3 in 1986. TLI provided a transport layer interface API. The ISO Transport Service Definition provided the model on which TLI is based. TLI provides an API between the OSI transport and session layers. TLI interfaces evolved further in AT&T System V, Release 4 version of UNIX and were also made available in SunOS 5.6 operating system interfaces.

XTI interfaces are an evolution of TLI interfaces and represent the future direction of this family of interfaces. Compatibility for applications using TLI interfaces is available. You do not need to port TLI applications to XTI immediately. New applications can use the XTI interfaces and you can port older applications to XTI when necessary.

TLI is implemented as a set of interface calls in a library (`libns1`) to which the applications link. XTI applications are compiled using the c89 front end and must be linked with the `xnet` library (`libxnet`). For additional information on compiling with XTI, see the [standards\(5\)](#) man page.

Note – An application using the XTI interface uses the `xti.h` header file, whereas an application using the TLI interface includes the `tiuser.h` header file.

XTI/TLI code can be independent of current transport providers when used in conjunction with some additional interfaces and mechanisms described in Chapter 4. The SunOS 5 product includes some transport providers (TCP, for example) as part of the base operating system. A transport provider performs services, and the transport user requests the services. The transport user issues service requests to the transport provider. An example is a request to transfer data over a connection TCP and UDP.

XTI/TLI can also be used for transport-independent programming by taking advantage of two components:

- Library routines that perform the transport services, in particular, transport selection and name-to-address translation. The network services library includes a set of interfaces that implement XTI/TLI for user processes. See [Chapter 10, “Transport Selection and Name-to-Address Mapping.”](#)
 Programs using TLI should be linked with the `libns1` network services library by specifying the `-l ns1` option at compile time.
 Programs using XTI should be linked with the `xnet` library by specifying the `-l xnet` option at compile time.
- State transition rules that define the sequence in which the transport routines can be invoked. For more information on state transition rules, see [“State Transitions” on page 204](#). The state tables define the legal sequence of library calls based on the state and the handling

of events. These events include user-generated library calls, as well as provider-generated event indications. XTI/TLI programmers should understand all state transitions before using the interface.

XTI/TLI Read/Write Interface

A user might want to establish a transport connection using `exec(2)` on an existing program (such as `/usr/bin/cat`) to process the data as it arrives over the connection. Existing programs use `read(2)` and `write(2)`. XTI/TLI does not directly support a read/write interface to a transport provider, but one is available. The interface enables you to issue `read(2)` and `write(2)` calls over a transport connection in the data transfer phase. This section describes the read/write interface to the connection mode service of XTI/TLI. This interface is not available with the connectionless mode service.

EXAMPLE 8-1 Read/Write Interface

```
#include <stropts.h>

/* Same local management and connection establishment steps. */

if (ioctl(fd, I_PUSH, "tirdwr") == -1) {
    perror("I_PUSH of tirdwr failed");
    exit(5);
}
close(0);
dup(fd);
execl("/usr/bin/cat", "/usr/bin/cat", (char *) 0);
perror("exec of /usr/bin/cat failed");
exit(6);
```

The client invokes the read/write interface by pushing `tirdwr` onto the stream associated with the transport endpoint. See the description of `I_PUSH` in the `streamio(7I)` man page. The `tirdwr` module converts XTI/TLI above the transport provider into a pure read/write interface. With the module in place, the client calls `close(2)` and `dup(2)` to establish the transport endpoint as its standard input file, and uses `/usr/bin/cat` to process the input.

Pushing `tirdwr` onto the transport provider forces XTI/TLI to use `read(2)` and `write(2)` semantics. XTI/TLI does not preserve message boundaries when using `read` and `write` semantics. Pop `tirdwr` from the transport provider to restore XTI/TLI semantics (see the description of `I_POP` in the `streamio(7I)` man page).



Caution – Push the `tirdwr` module onto a stream only when the transport endpoint is in the data transfer phase. After pushing the module, the user cannot call any XTI/TLI routines. If the user invokes an XTI/TLI routine, `tirdwr` generates a fatal protocol error, `EPROTO`, on the stream, rendering it unusable. If you then pop the `tirdwr` module off the stream, the transport connection aborts. See the description of `I_POP` in the [streamio\(7I\)](#) man page.

Write Data

After you send data over the transport connection with `write(2)`, `tirdwr` passes data through to the transport provider. If you send a zero-length data packet, which the mechanism allows, `tirdwr` discards the message. If the transport connection is aborted, a hang-up condition is generated on the stream, further `write(2)` calls fail, and `errno` is set to `ENXIO`. This problem might occur, for example, because the remote user aborts the connection using `t_snddis(3NSL)`. You can still retrieve any available data after a hang-up.

Read Data

Receive data that arrives at the transport connection with `read(2)`. `tirdwr` passes data from the transport provider. The `tirdwr` module processes any other event or request passed to the user from the provider as follows:

- `read(2)` cannot identify expedited data to the user. If `read(2)` receives an expedited data request, `tirdwr` generates a fatal protocol error, `EPROTO`, on the stream. The error causes further system calls to fail. Do not use `read(2)` to receive expedited data.
- `tirdwr` discards an abortive disconnect request and generates a hang-up condition on the stream. Subsequent `read(2)` calls retrieve any remaining data, then return zero for all further calls, indicating end of file.
- `tirdwr` discards an orderly release request and delivers a zero-length message to the user. As described in the `read(2)` man page, this notifies the user of end of file by returning 0.
- If `read(2)` receives any other XTI/TLI request, `tirdwr` generates a fatal protocol error, `EPROTO`, on the stream. This causes further system calls to fail. If a user pushes `tirdwr` onto a stream after establishing the connection, `tirdwr` generates no request.

Close Connection

With `tirdwr` on a stream, you can send and receive data over a transport connection for the duration of the connection. Either user can terminate the connection by closing the file descriptor associated with the transport endpoint or by popping the `tirdwr` module off the stream. In either case, `tirdwr` does the following:

- If `tirdwr` receives an orderly release request, it passes the request to the transport provider to complete the orderly release of the connection. The remote user who initiated the orderly release procedure receives the expected request when data transfer completes.
- If `tirdwr` receives a disconnect request, it takes no special action.
- If `tirdwr` receives neither an orderly release nor a disconnect request, it passes a disconnect request to the transport provider to abort the connection.
- If an error occurs on the stream and `tirdwr` does not receive a disconnect request, it passes a disconnect request to the transport provider.

A process cannot initiate an orderly release after pushing `tirdwr` onto a stream. `tirdwr` handles an orderly release if the user on the other side of a transport connection initiates the release. If the client in this section is communicating with a server program, the server terminates the transfer of data with an orderly release request. The server then waits for the corresponding request from the client. At that point, the client exits and closes the transport endpoint. After closing the file descriptor, `tirdwr` initiates the orderly release request from the client's side of the connection. This release generates the request on which the server blocks.

Some protocols, like TCP, require this orderly release to ensure intact delivery of the data.

Advanced XTI/TLI Topics

This section presents additional XTI/TLI concepts:

- [“Asynchronous Execution Mode” on page 193](#) describes optional nonblocking (asynchronous) mode for some library calls.
- [“Advanced XTI/TLI Programming Example” on page 194](#) is a program example of a server supporting multiple outstanding connect requests and operating in an event-driven manner.

Asynchronous Execution Mode

Many XTI/TLI library routines block to wait for an incoming event. However, some time-critical applications should not block for any reason. An application can do local processing while waiting for some asynchronous XTI/TLI event.

Applications can access asynchronous processing of XTI/TLI events through the combination of asynchronous features and the non-blocking mode of XTI/TLI library routines. See the *ONC+ Developer's Guide* for information on use of the `poll(2)` system call and the `I_SETSIG` `ioctl(2)` command to process events asynchronously.

You can run each XTI/TLI routine that blocks for an event in a special non-blocking mode. For example, `t_listen(3NSL)` normally blocks for a connect request. A server can periodically poll a transport endpoint for queued connect requests by calling `t_listen(3NSL)` in the

non-blocking (or asynchronous) mode. You enable the asynchronous mode by setting `O_NDELAY` or `O_NONBLOCK` in the file descriptor. Set these modes as a flag through `t_open(3NSL)`, or by calling `fcntl(2)` before calling the XTI/TLI routine. Use `fcntl(2)` to enable or disable this mode at any time. All program examples in this chapter use the default synchronous processing mode.

Use of `O_NDELAY` or `O_NONBLOCK` affects each XTI/TLI routine differently. You need to determine the exact semantics of `O_NDELAY` or `O_NONBLOCK` for a particular routine.

Advanced XTI/TLI Programming Example

[Example 8–2](#) demonstrates two important concepts. The first is a server's ability to manage multiple outstanding connect requests. The second is event-driven use of XTI/TLI and the system call interface.

By using XTI/TLI, a server can manage multiple outstanding connect requests. One reason to receive several simultaneous connect requests is to prioritize the clients. A server can receive several connect requests, and accept them in an order based on the priority of each client.

The second reason for handling several outstanding connect requests is to overcome the limits of single-threaded processing. Depending on the transport provider, while a server is processing one connect request, other clients see the server as busy. If multiple connect requests are processed simultaneously, the server is busy only if more than the maximum number of clients try to call the server simultaneously.

The server example is event-driven: the process polls a transport endpoint for incoming XTI/TLI events and takes the appropriate actions for the event received. The example following demonstrates the ability to poll multiple transport endpoints for incoming events.

EXAMPLE 8–2 Endpoint Establishment (Convertible to Multiple Connections)

```
#include <tiuser.h>
#include <fcntl.h>
#include <stdio.h>
#include <poll.h>
#include <stropts.h>
#include <signal.h>

#define NUM_FDS 1
#define MAX_CONN_IND 4
#define SRV_ADDR 1 /* server's well known address */

int conn_fd; /* server connection here */
extern int t_errno;
/* holds connect requests */
struct t_call *calls[NUM_FDS][MAX_CONN_IND];

main()
{
```

EXAMPLE 8-2 Endpoint Establishment (Convertible to Multiple Connections) *(Continued)*

```

struct pollfd pollfds[NUM_FDS];
struct t_bind *bind;
int i;

/*
 * Only opening and binding one transport endpoint, but more can
 * be supported
 */
if ((pollfds[0].fd = t_open("/dev/tivc", O_RDWR,
    (struct t_info *) NULL)) == -1) {
    t_error("t_open failed");
    exit(1);
}
if ((bind = (struct t_bind *) t_alloc(pollfds[0].fd, T_BIND,
    T_ALL)) == (struct t_bind *) NULL) {
    t_error("t_alloc of t_bind structure failed");
    exit(2);
}
bind->qlen = MAX_CONN_IND;
bind->addr.len = sizeof(int);
*(int *) bind->addr.buf = SRV_ADDR;
if (t_bind(pollfds[0].fd, bind, bind) == -1) {
    t_error("t_bind failed");
    exit(3);
}
/* Was the correct address bound? */
if (bind->addr.len != sizeof(int) ||
    *(int *)bind->addr.buf != SRV_ADDR) {
    fprintf(stderr, "t_bind bound wrong address\n");
    exit(4);
}
}

```

The file descriptor returned by `t_open(3NSL)` is stored in a `pollfd` structure that controls polling of the transport endpoints for incoming data. See the `poll(2)` man page. Only one transport endpoint is established in this example. However, the remainder of the example is written to manage multiple transport endpoints. Several endpoints could be supported with minor changes to [Example 8-2](#).

This server sets `qlen` to a value greater than 1 for `t_bind(3NSL)`. This value specifies that the server should queue multiple outstanding connect requests. The server accepts the current connect request before accepting additional connect requests. This example can queue up to `MAX_CONN_IND` connect requests. The transport provider can negotiate the value of `qlen` to be smaller if the provider cannot support `MAX_CONN_IND` outstanding connect requests.

After the server binds its address and is ready to process connect requests, it behaves as shown in the following example.

EXAMPLE 8-3 Processing Connection Requests

```
pollfds[0].events = POLLIN;
```

EXAMPLE 8-3 Processing Connection Requests (Continued)

```

while (TRUE) {
    if (poll(pollfds, NUM_FDS, -1) == -1) {
        perror("poll failed");
        exit(5);
    }
    for (i = 0; i < NUM_FDS; i++) {
        switch (pollfds[i].revents) {
            default:
                perror("poll returned error event");
                exit(6);
            case 0:
                continue;
            case POLLIN:
                do_event(i, pollfds[i].fd);
                service_conn_ind(i, pollfds[i].fd);
        }
    }
}

```

The `revents` field of the `pollfd` structure is set to `POLLIN`, which notifies the server of any incoming XTI/TLI events. The server then enters an infinite loop in which it polls the transport endpoints for events, and processes events as they occur.

The `poll(2)` call blocks indefinitely for an incoming event. On return, the server checks the value of `revents` for each entry, one per transport endpoint, for new events. If `revents` is `0`, the endpoint has generated no events and the server continues to the next endpoint. If `revents` is `POLLIN`, there is an event on the endpoint. The server calls `do_event` to process the event. Any other value in `revents` indicates an error on the endpoint, and the server exits. With multiple endpoints, the server should close this descriptor and continue.

Each time the server iterates the loop, it calls `service_conn_ind` to process any outstanding connect requests. If another connect request is pending, `service_conn_ind` saves the new connect request and responds to it later.

The server calls `do_event` in the following example to process an incoming event.

EXAMPLE 8-4 Event Processing Routine

```

do_event( slot, fd)
int slot;
int fd;
{
    struct t_discon *discon;
    int i;

    switch (t_look(fd)) {
        default:
            fprintf(stderr, "t_look: unexpected event\n");
            exit(7);
        case T_ERROR:

```

EXAMPLE 8-4 Event Processing Routine (Continued)

```

        fprintf(stderr, "t_look returned T_ERROR event\n");
        exit(8);
    case -1:
        t_error("t_look failed");
        exit(9);
    case 0:
        /* since POLLIN returned, this should not happen */
        fprintf(stderr, "t_look returned no event\n");
        exit(10);
    case T_LISTEN:
        /* find free element in calls array */
        for (i = 0; i < MAX_CONN_IND; i++) {
            if (calls[slot][i] == (struct t_call *) NULL)
                break;
        }
        if ((calls[slot][i] = (struct t_call *) t_alloc( fd, T_CALL,
            T_ALL)) == (struct t_call *) NULL) {
            t_error("t_alloc of t_call structure failed");
            exit(11);
        }
        if (t_listen(fd, calls[slot][i] ) == -1) {
            t_error("t_listen failed");
            exit(12);
        }
        break;
    case T_DISCONNECT:
        discon = (struct t_discon *) t_alloc(fd, T_DIS, T_ALL);
        if (discon == (struct t_discon *) NULL) {
            t_error("t_alloc of t_discon structure failed");
            exit(13)
        }
        if(t_rcvdis( fd, discon) == -1) {
            t_error("t_rcvdis failed");
            exit(14);
        }
        /* find call ind in array and delete it */
        for (i = 0; i < MAX_CONN_IND; i++) {
            if (discon->sequence == calls[slot][i]->sequence) {
                t_free(calls[slot][i], T_CALL);
                calls[slot][i] = (struct t_call *) NULL;
            }
        }
        t_free(discon, T_DIS);
        break;
    }
}

```

The arguments in [Example 8-4](#) are a number (*slot*) and a file descriptor (*fd*). A *slot* is the index into the global array `calls`, which has an entry for each transport endpoint. Each entry is an array of `t_call` structures that hold incoming connect requests for the endpoint.

The `do_event` module calls `t_look(3NSL)` to identify the XTI/TLI event on the endpoint specified by `fd`. If the event is a connect request (`T_LISTEN` event) or disconnect request (`T_DISCONNECT` event), the event is processed. Otherwise, the server prints an error message and exits.

For connect requests, `do_event` scans the array of outstanding connect requests for the first free entry. A `t_call` structure is allocated for the entry, and the connect request is received by `t_listen(3NSL)`. The array is large enough to hold the maximum number of outstanding connect requests. The processing of the connect request is deferred.

A disconnect request must correspond to an earlier connect request. The `do_event` module allocates a `t_discon` structure to receive the request. This structure has the following fields:

```
struct t_discon {
    struct netbuf  udata;
    int           reason;
    int           sequence;
}
```

The `udata` structure contains any user data sent with the disconnect request. The value of `reason` contains a protocol-specific disconnect reason code. The value of `sequence` identifies the connect request that matches the disconnect request.

The server calls `t_rcvdis(3NSL)` to receive the disconnect request. The array of connect requests is scanned for one that contains the sequence number that matches the sequence number in the disconnect request. When the connect request is found, its structure is freed and the entry is set to `NULL`.

When an event is found on a transport endpoint, `service_conn_ind` is called to process all queued connect requests on the endpoint, as the following example shows.

EXAMPLE 8-5 Process All Connect Requests

```
service_conn_ind(slot, fd)
{
    int i;

    for (i = 0; i < MAX_CONN_IND; i++) {
        if (calls[slot][i] == (struct t_call *) NULL)
            continue;
        if ((conn_fd = t_open( "/dev/tivc", 0_RDWR,
            (struct t_info *) NULL)) == -1) {
            t_error("open failed");
            exit(15);
        }
        if (t_bind(conn_fd, (struct t_bind *) NULL,
            (struct t_bind *) NULL) == -1) {
            t_error("t_bind failed");
            exit(16);
        }
        if (t_accept(fd, conn_fd, calls[slot][i]) == -1) {
            if (t_errno == TLOOK) {
```

EXAMPLE 8-5 Process All Connect Requests (Continued)

```

        t_close(conn_fd);
        return;
    }
    t_error("t_accept failed");
    exit(167);
}
t_free(calls[slot][i], T_CALL);
calls[slot][i] = (struct t_call *) NULL;
run_server(fd);
}
}

```

For each transport endpoint, the array of outstanding connect requests is scanned. For each request, the server opens a responding transport endpoint, binds an address to the endpoint, and accepts the connection on the endpoint. If another connect or disconnect request arrives before the current request is accepted, `t_accept(3NSL)` fails and sets `t_errno` to `TL00K`. You cannot accept an outstanding connect request if any pending connect request events or disconnect request events exist on the transport endpoint.

If this error occurs, the responding transport endpoint is closed and `service_conn_ind` returns immediately, saving the current connect request for later processing. This activity causes the server's main processing loop to be entered, and the new event is discovered by the next call to `poll(2)`. In this way, the user can queue multiple connect requests.

Eventually, all events are processed, and `service_conn_ind` is able to accept each connect request in turn.

Asynchronous Networking

This section discusses the techniques of asynchronous network communication using XTI/TLI for real-time applications. The SunOS platform provides support for asynchronous network processing of XTI/TLI events using a combination of STREAMS asynchronous features and the non-blocking mode of the XTI/TLI library routines.

Networking Programming Models

Like file and device I/O, network transfers can be done synchronously or asynchronously with process service requests.

Synchronous networking proceeds similar to synchronous file and device I/O. Like the `write(2)` interface, the send request returns after buffering the message, but might suspend the calling process if buffer space is not immediately available. Like the `read(2)` interface, a receive request suspends execution of the calling process until data arrives to satisfy the request.

Because there are no guaranteed bounds for transport services, synchronous networking is inappropriate for processes that must have real-time behavior with respect to other devices.

Asynchronous networking is provided by non-blocking service requests. Additionally, applications can request asynchronous notification when a connection might be established, when data might be sent, or when data might be received.

Asynchronous Connectionless-Mode Service

Asynchronous connectionless mode networking is conducted by configuring the endpoint for non-blocking service, and either polling for or receiving asynchronous notification when data might be transferred. If asynchronous notification is used, the actual receipt of data typically takes place within a signal handler.

Making the Endpoint Asynchronous

After the endpoint has been established using `t_open(3NSL)`, and its identity established using `t_bind(3NSL)`, the endpoint can be configured for asynchronous service. Use the `fcntl(2)` interface to set the `O_NONBLOCK` flag on the endpoint. Thereafter, calls to `t_sndudata(3NSL)` for which no buffer space is immediately available return -1 with `t_errno` set to `TFLOW`. Likewise, calls to `t_rcvudata(3NSL)` for which no data are available return -1 with `t_errno` set to `TNODATA`.

Asynchronous Network Transfers

Although an application can use `poll(2)` to check periodically for the arrival of data or to wait for the receipt of data on an endpoint, receiving asynchronous notification when data arrives might be necessary. Use `ioctl(2)` with the `I_SETSIG` command to request that a `SIGPOLL` signal be sent to the process upon receipt of data at the endpoint. Applications should check for the possibility of multiple messages causing a single signal.

In the following example, `protocol` is the name of the application-chosen transport protocol.

```
#include <sys/types.h>
#include <tiuser.h>
#include <signal.h>
#include <stropts.h>

int          fd;
struct t_bind *bind;
void        sigpoll(int);

    fd = t_open(protocol, O_RDWR, (struct t_info *) NULL);

    bind = (struct t_bind *) t_alloc(fd, T_BIND, T_ADDR);
    ... /* set up binding address */
    t_bind(fd, bind, bin
```

```

/* make endpoint non-blocking */
fcntl(fd, F_SETFL, fcntl(fd, F_GETFL) | O_NONBLOCK);

/* establish signal handler for SIGPOLL */
signal(SIGPOLL, sigpoll);

/* request SIGPOLL signal when receive data is available */
ioctl(fd, I_SETSIG, S_INPUT | S_HIPRI);

...

void sigpoll(int sig)
{
    int                flags;
    struct t_unitdata  ud;

    for (;;) {
        ... /* initialize ud */
        if (t_rcvudata(fd, &ud, &flags) < 0) {
            if (t_errno == TNODATA)
                break; /* no more messages */
            ... /* process other error conditions */
        }
        ... /* process message in ud */
    }
}

```

Asynchronous Connection-Mode Service

For connection-mode service, an application can arrange not only for the data transfer, but also for the establishment of the connection itself to be done asynchronously. The sequence of operations depends on whether the process is attempting to connect to another process or is awaiting connection attempts.

Asynchronously Establishing a Connection

A process can attempt a connection and asynchronously complete the connection. The process first creates the connecting endpoint and, using `fcntl(2)`, configures the endpoint for non-blocking operation. As with connectionless data transfers, the endpoint can also be configured for asynchronous notification upon completion of the connection and subsequent data transfers. The connecting process then uses `t_connect(3NSL)` to initiate setting up the transfer. Then `t_rcvconnect(3NSL)` is used to confirm the establishment of the connection.

Asynchronous Use of a Connection

To asynchronously await connections, a process first establishes a non-blocking endpoint bound to a service address. When either the result of `poll(2)` or an asynchronous notification indicates that a connection request has arrived, the process can get the connection request by using `t_listen(3NSL)`. To accept the connection, the process uses `t_accept(3NSL)`. The responding endpoint must be separately configured for asynchronous data transfers.

The following example illustrates how to request a connection asynchronously.

```
#include <tiuser.h>
int      fd;
struct t_call  *call;

fd = /* establish a non-blocking endpoint */

call = (struct t_call *) t_alloc(fd, T_CALL, T_ADDR);
/* initialize call structure */
t_connect(fd, call, call);

/* connection request is now proceeding asynchronously */

/* receive indication that connection has been accepted */
t_rcvconnect(fd, &call);
```

The following example illustrates listening for connections asynchronously.

```
#include <tiuser.h>
int      fd, res_fd;
struct t_call  call;

fd = /* establish non-blocking endpoint */

/*receive indication that connection request has arrived */
call = (struct t_call *) t_alloc(fd, T_CALL, T_ALL);
t_listen(fd, &call);

/* determine whether or not to accept connection */
res_fd = /* establish non-blocking endpoint for response */
t_accept(fd, res_fd, call);
```

Asynchronous Open

Occasionally, an application might be required to dynamically open a regular file in a file system mounted from a remote host, or on a device whose initialization might be prolonged. However, while such a request to open a file is being processed, the application is unable to achieve real-time response to other events. The SunOS software solves this problem by having a second process handle the actual opening of the file, then passes the file descriptor to the real-time process.

Transferring a File Descriptor

The STREAMS interface provided by the SunOS platform provides a mechanism for passing an open file descriptor from one process to another. The process with the open file descriptor uses [ioctl\(2\)](#) with a command argument of `I_SENDFD`. The second process obtains the file descriptor by calling [ioctl\(2\)](#) with a command argument of `I_RECVFD`.

In the following example, the parent process prints out information about the test file, and creates a pipe. Next, the parent creates a child process that opens the test file and passes the open file descriptor back to the parent through the pipe. The parent process then displays the status information on the new file descriptor.

EXAMPLE 8-6 File Descriptor Transfer

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <stropts.h>
#include <stdio.h>

#define TESTFILE "/dev/null"
main(int argc, char *argv[])
{
    int fd;
    int pipefd[2];
    struct stat statbuf;

    stat(TESTFILE, &statbuf);
    statout(TESTFILE, &statbuf);
    pipe(pipefd);
    if (fork() == 0) {
        close(pipefd[0]);
        sendfd(pipefd[1]);
    } else {
        close(pipefd[1])
        recvfd(pipefd[0]);
    }
}

sendfd(int p)
{
    int tfd;

    tfd = open(TESTFILE, O_RDWR);
    ioctl(p, I_SENDFD, tfd);
}

recvfd(int p)
{
    struct strrecvfd rdbuf;
    struct stat statbuf;
    char          fdbuf[32];

    ioctl(p, I_RECVFD, &rdbuf);
    fstat(rdbuf.fd, &statbuf);
    sprintf(fdbuf, "recvfd=%d", rdbuf.fd);
    statout(fdbuf, &statbuf);
}

statout(char *f, struct stat *s)
{
    printf("stat: from=%s mode=0%o, ino=%ld, dev=%lx, rdev=%lx\n",
        f, s->st_mode, s->st_ino, s->st_dev, s->st_rdev);
    fflush(stdout);
}
```

EXAMPLE 8-6 File Descriptor Transfer (Continued)

```
}
```

State Transitions

The tables in the following sections describe all state transitions associated with XTI/TLI.

XTI/TLI States

The following table defines the states used in XTI/TLI state transitions, along with the service types.

TABLE 8-1 XTI/TLI State Transitions and Service Types

State	Description	Service Type
T_UNINIT	Uninitialized—initial and final state of interface	T_COTS, T_COTS_ORD, T_CLTS
T_UNBND	Initialized but not bound	T_COTS, T_COTS_ORD, T_CLTS
T_IDLE	No connection established	T_COTS, T_COTS_ORD, T_CLTS
T_OUTCON	Outgoing connection pending for client	T_COTS, T_COTS_ORD
T_INCON	Incoming connection pending for server	T_COTS, T_COTS_ORD
T_DATAXFER	Data transfer	T_COTS, T_COTS_ORD
T_OUTREL	Outgoing orderly release (waiting for orderly release request)	T_COTS_ORD
T_INREL	Incoming orderly release (waiting to send orderly release request)	T_COTS_ORD

Outgoing Events

The outgoing events described in the following table correspond to the status returned from the specified transport routines, where these routines send a request or response to the transport provider. In the table, some events, such as “accept,” are distinguished by the context in which they occur. The context is based on the values of the following variables:

- *ocnt* – Count of outstanding connect requests
- *fd* – File descriptor of the current transport endpoint
- *resfd* – File descriptor of the transport endpoint where a connection is accepted

TABLE 8-2 Outgoing Events

Event	Description	Service Type
opened	Successful return of <code>t_open(3NSL)</code>	T_COTS, T_COTS_ORD, T_CLTS
bind	Successful return of <code>t_bind(3NSL)</code>	T_COTS, T_COTS_ORD, T_CLTS
optmgmt	Successful return of <code>t_optmgmt(3NSL)</code>	T_COTS, T_COTS_ORD, T_CLTS
unbind	Successful return of <code>t_unbind(3NSL)</code>	T_COTS, T_COTS_ORD, T_CLTS
closed	Successful return of <code>t_close(3NSL)</code>	T_COTS, T_COTS_ORD, T_CLT
connect1	Successful return of <code>t_connect(3NSL)</code> in synchronous mode	T_COTS, T_COTS_ORD
connect2	TNODATA error on <code>t_connect(3NSL)</code> in asynchronous mode, or TLOOK error due to a disconnect request arriving on the transport endpoint	T_COTS, T_COTS_ORD
accept1	Successful return of <code>t_accept(3NSL)</code> with <code>ocnt == 1, fd == resfd</code>	T_COTS, T_COTS_ORD
accept2	Successful return of <code>t_accept(3NSL)</code> with <code>ocnt == 1, fd != resfd</code>	T_COTS, T_COTS_ORD
accept3	Successful return of <code>t_accept(3NSL)</code> with <code>ocnt > 1</code>	T_COTS, T_COTS_ORD
snd	Successful return of <code>t_snd(3NSL)</code>	T_COTS, T_COTS_ORD
snddis1	Successful return of <code>t_snddis(3NSL)</code> with <code>ocnt <= 1</code>	T_COTS, T_COTS_ORD
snddis2	Successful return of <code>t_snddis(3NSL)</code> with <code>ocnt > 1</code>	T_COTS, T_COTS_ORD
sndrel	Successful return of <code>t_sndrel(3NSL)</code>	T_COTS_ORD
sndudata	Successful return of <code>t_sndudata(3NSL)</code>	T_CLTS

Incoming Events

The incoming events correspond to the successful return of the specified routines. These routines return data or event information from the transport provider. The only incoming event not associated directly with the return of a routine is `pass_conn`, which occurs when a connection is transferred to another endpoint. The event occurs on the endpoint that is being passed the connection, although no XTI/TLI routine is called on the endpoint.

In the following table, the `rcvdis` events are distinguished by the value of `ocnt`, the count of outstanding connect requests on the endpoint.

TABLE 8-3 Incoming Events

Event	Description	Service Type
listen	Successful return of <code>t_listen(3NSL)</code>	T_COTS, T_COTS_ORD
rcvconnect	Successful return of <code>t_rcvconnect(3NSL)</code>	T_COTS, T_COTS_ORD
rcv	Successful return of <code>t_rcv(3NSL)</code>	T_COTS, T_COTS_ORD
rcvdis1	Successful return of <code>t_rcvdis(3NSL)</code> <code>rcvdis1t_rcvdis()</code> , <code>onct <= 0</code>	T_COTS, T_COTS_ORD
rcvdis2	Successful return of <code>t_rcvdis(3NSL)</code> , <code>onct == 1</code>	T_COTS, T_COTS_ORD
rcvdis3	Successful return of <code>t_rcvdis(3NSL)</code> with <code>onct > 1</code>	T_COTS, T_COTS_ORD
rcvrel	Successful return of <code>t_rcvrel(3NSL)</code>	T_COTS_ORD
rcvudata	Successful return of <code>t_rcvudata(3NSL)</code>	T_CLTS
rcvuderr	Successful return of <code>t_rcvuderr(3NSL)</code>	T_CLTS
pass_conn	Receive a passed connection	T_COTS, T_COTS_ORD

State Tables

The state tables describe the XTI/TLI state transitions. Each box contains the next state, given the current state (column) and the current event (row). An empty box is an invalid state/event combination. Each box can also have an action list. Actions must be done in the order specified in the box.

You should understand the following when studying the state tables:

- `t_close(3NSL)` terminates an established connection for a connection-oriented transport provider. The connection termination will be either orderly or abortive, depending on the service type supported by the transport provider. See the `t_getinfo(3NSL)` man page.
- If a transport user issues a interface call out of sequence, the interface fails and `t_errno` is set to TOUTSTATE. The state does not change.
- The error codes TLOOK or TNODATA after `t_connect(3NSL)` can result in state changes. The state tables assume correct use of XTI/TLI.
- Any other transport error does not change the state, unless the man page for the interface says otherwise.
- The support interfaces `t_getinfo(3NSL)`, `t_getstate(3NSL)`, `t_alloc(3NSL)`, `t_free(3NSL)`, `t_sync(3NSL)`, `t_look(3NSL)`, and `t_error(3NSL)` are excluded from the state tables because they do not affect the state.

Some of the state transitions listed in the tables below offer actions the transport user must take. Each action is represented by a digit derived from the list below:

- Set the count of outstanding connect requests to zero
- Increment the count of outstanding connect requests
- Decrement the count of outstanding connect requests
- Pass a connection to another transport endpoint, as indicated in the `t_accept(3NSL)` man page

The following table shows endpoint establishment states.

TABLE 8-4 Connection Establishment State

Event/State	T_UNINIT	T_UNBND	T_IDLE
opened	T_UNBND		
bind		T_IDLE [1]	
optmgmt (TLI only)			T_IDLE
unbind			T_UNBND
closed		T_UNINIT	

The following table shows data transfer in connection mode.

TABLE 8-5 Connection Mode State: Part 1

Event/State	T_IDLE	T_OUTCON	T_INCON	T_DATAXFER
connect1	T_DATAXFER			
connect2	T_OUTCON			
rcvconnect		T_DATAXFER		
listen	T_INCON [2]		T_INCON [2]	
accept1			T_DATAXFER [3]	
accept2			T_IDLE [3] [4]	
accept3			T_INCON [3] [4]	
snd				T_DATAXFER
rcv				T_DATAXFER
snddis1		T_IDLE	T_IDLE [3]	T_IDLE
snddis2			T_INCON [3]	

TABLE 8-5 Connection Mode State: Part 1 (Continued)

Event/State	T_IDLE	T_OUTCON	T_INCON	T_DATAXFER
rcvdis1		T_IDLE		T_IDLE
rcvdis2			T_IDLE [3]	
rcvdis3			T_INCON [3]	
sndrel				T_OUTREL
rcvrel				T_INREL
pass_conn	T_DATAXFER			
optmgmt	T_IDLE	T_OUTCON	T_INCON	T_DATAXFER
closed	T_UNINIT	T_UNINIT	T_UNINIT	T_UNINIT

The following table shows connection establishment/connection release/data transfer in connection mode.

TABLE 8-6 Connection Mode State: Part 2

Event/State	T_OUTREL	T_INREL	T_UNBND
connect1			
connect2			
rcvconnect			
listen			
accept1			
accept2			
accept3			
snd		T_INREL	
rcv	T_OUTREL		
snddis1	T_IDLE	T_IDLE	
snddis2			
rcvdis1	T_IDLE	T_IDLE	
rcvdis2			
rcvdis3			
sndrel		T_IDLE	

TABLE 8-6 Connection Mode State: Part 2 (Continued)

Event/State	T_OUTREL	T_INREL	T_UNBND
rcvrel	T_IDLE		
pass_conn			T_DATAXFER
optmgmt	T_OUTREL	T_INREL	T_UNBND
closed	T_UNINIT	T_UNINIT	

The following table shows connectionless mode states.

TABLE 8-7 Connectionless Mode State

Event/State	T_IDLE
snudata	T_IDLE
rcvdata	T_IDLE
rcvuderr	T_IDLE

Guidelines to Protocol Independence

The set of XTI/TLI services, common to many transport protocols, offers protocol independence to applications. Not all transport protocols support all XTI/TLI services. If software must run in a variety of protocol environments, use only the common services.

The following is a list of services that might not be common to all transport protocols.

- In connection mode service, a transport service data unit (TSDU) might not be supported by all transport providers. Make no assumptions about preserving logical data boundaries across a connection.
- Protocol and implementation-specific service limits are returned by the `t_open(3NSL)` and `t_getinfo(3NSL)` routines. Use these limits to allocate buffers to store protocol-specific transport addresses and options.
- Do not send user data with connect requests or disconnect requests, such as `t_connect(3NSL)` and `t_snddis(3NSL)`. Not all transport protocols can use this method.
- The buffers in the `t_call` structure used for `t_listen(3NSL)` must be large enough to hold any data sent by the client during connection establishment. Use the `T_ALL` argument to `t_alloc(3NSL)` to set maximum buffer sizes to store the address, options, and user data for the current transport provider.

- Do not specify a protocol address on `t_bind(3NSL)` on a client-side endpoint. The transport provider should assign an appropriate address to the transport endpoint. A server should retrieve its address for `t_bind(3NSL)` in a way that does not require knowledge of the transport provider's name space.
- Do not make assumptions about formats of transport addresses. Transport addresses should not be constants in a program. [Chapter 10, “Transport Selection and Name-to-Address Mapping,”](#) contains detailed information about transport selection.
- The reason codes associated with `t_rcvdis(3NSL)` are protocol-dependent. Do not interpret these reason codes if protocol independence is important.
- The `t_rcvuderr(3NSL)` error codes are protocol dependent. Do not interpret these error codes if protocol independence is a concern.
- Do not code the names of devices into programs. The device node identifies a particular transport provider and is not protocol independent. See [Chapter 10, “Transport Selection and Name-to-Address Mapping,”](#) for details regarding transport selection.
- Do not use the optional orderly release facility of the connection mode service, provided by `t_sndrel(3NSL)` and `t_rcvrel(3NSL)`, in programs targeted for multiple protocol environments. This facility is not supported by all connection-based transport protocols. Using the facility can prevent programs from successfully communicating with open systems.

XTI/TLI Versus Socket Interfaces

XTI/TLI and sockets are different methods of handling the same tasks. Although they provide mechanisms and services that are functionally similar, they do not provide one-to-one compatibility of routines or low-level services. Observe the similarities and differences between the XTI/TLI and socket-based interfaces before you decide to port an application.

The following issues are related to transport independence, and can have some bearing on RPC applications:

- Privileged ports – Privileged ports are an artifact of the Berkeley Software Distribution (BSD) implementation of the TCP/IP Internet Protocols. These ports are not portable. The notion of privileged ports is not supported in the transport-independent environment.
- Opaque addresses – Separating the portion of an address that names a host from the portion of an address that names the service at that host cannot be done in a transport-independent fashion. Be sure to change any code that assumes it can discern the host address of a network service.
- Broadcast – No transport-independent form of broadcast address exists.

Socket-to-XTI/TLI Equivalents

The following table shows approximate equivalents between XTI/TLI interfaces and socket interfaces. The comment field describes the differences. If the comment column is blank, either the interfaces are similar or no equivalent interface exists in either interface.

TABLE 8-8 TLI and Socket Equivalent Functions

TLI interface	Socket interface	Comments
<code>t_open(3NSL)</code>	<code>socket(3SOCKET)</code>	
–	<code>socketpair(3SOCKET)</code>	
<code>t_bind(3NSL)</code>	<code>bind(3SOCKET)</code>	<code>t_bind(3NSL)</code> sets the queue depth for passive sockets, but <code>bind(3SOCKET)</code> does not. For sockets, the queue length is specified in the call to <code>listen(3SOCKET)</code> .
<code>t_optmgmt(3NSL)</code>	<code>getsockopt(3SOCKET)</code> <code>setsockopt(3SOCKET)</code>	<code>t_optmgmt(3NSL)</code> manages only transport options. <code>getsockopt(3SOCKET)</code> and <code>setsockopt(3SOCKET)</code> can manage options at the transport layer, but also at the socket layer and at the arbitrary protocol layer.
<code>t_unbind(3NSL)</code>	–	
<code>t_close(3NSL)</code>	<code>close(2)</code>	
<code>t_getinfo(3NSL)</code>	<code>getsockopt(3SOCKET)</code>	<code>t_getinfo(3NSL)</code> returns information about the transport. <code>getsockopt(3SOCKET)</code> can return information about the transport and the socket.
<code>t_getstate(3NSL)</code>	–	
<code>t_sync(3NSL)</code>	–	
<code>t_alloc(3NSL)</code>	–	
<code>t_free(3NSL)</code>	–	
<code>t_look(3NSL)</code>	–	<code>getsockopt(3SOCKET)</code> with the <code>SO_ERROR</code> option returns the same kind of error information as <code>t_look(3NSL)t_look()</code> .
<code>t_error(3NSL)</code>	<code>perror(3C)</code>	

TLI interface	Socket interface	Comments
<code>t_connect(3NSL)</code>	<code>connect(3SOCKET)</code>	You do not need to bind the local endpoint before invoking <code>connect(3SOCKET)</code> . Bind the endpoint before calling <code>t_connect(3NSL)</code> . You can use <code>connect(3SOCKET)</code> on a connectionless endpoint to set the default destination address for datagrams. You can send data using <code>connect(3SOCKET)</code> .
<code>t_rcvconnect(3NSL)</code>	-	
<code>t_listen(3NSL)</code>	<code>listen(3SOCKET)</code>	<code>t_listen(3NSL)</code> waits for connection indications. <code>listen(3SOCKET)</code> sets the queue depth.
<code>t_accept(3NSL)</code>	<code>accept(3SOCKET)</code>	
<code>t_snd(3NSL)</code>	<code>send(3SOCKET)</code> <code>sendto(3SOCKET)</code> <code>sendmsg(3SOCKET)</code>	<code>sendto(3SOCKET)</code> and <code>sendmsg(3SOCKET)</code> operate in connection mode as well as in datagram mode.
<code>t_rcv(3NSL)</code>	<code>recv(3SOCKET)</code> <code>recvfrom(3SOCKET)</code> <code>recvmsg(3SOCKET)</code>	<code>recvfrom(3SOCKET)</code> and <code>recvmsg(3SOCKET)</code> operate in connection mode as well as datagram mode.
<code>t_snddis(3NSL)</code>	-	
<code>t_rcvdis(3NSL)</code>	-	
<code>t_sndrel(3NSL)</code>	<code>shutdown(3SOCKET)</code>	
<code>t_rcvrel(3NSL)</code>	-	
<code>t_sndudata(3NSL)</code>	<code>sendto(3SOCKET)</code> <code>recvmsg(3SOCKET)</code>	
<code>t_rcvuderr(3NSL)</code>	-	
<code>read(2), write(2)</code>	<code>read(2), write(2)</code>	In XTI/TLI you must push the <code>tirdwr(7M)</code> module before calling <code>read(2)</code> or <code>write(2)</code> . In sockets, calling <code>read(2)</code> or <code>write(2)</code> suffices.

Additions to the XTI Interface

The XNS 5 (UNIX03) standard introduces some new XTI interfaces. These are briefly described below. You can find the details in the relevant manual pages. These interfaces are not available for TLI users. The scatter-gather data transfer interfaces are:

<code>t_sndvudata(3NSL)</code>	Send a data unit from one or more non-contiguous buffers
<code>t_rcvvudata(3NSL)</code>	Receive a data unit into one or more non-contiguous buffers
<code>t_sndv(3NSL)</code>	Send data or expedited data from one or more non-contiguous buffers on a connection
<code>t_rcvv(3NSL)</code>	Receive data or expedited data sent over a connection and put the data into one or more non-contiguous buffers

The XTI utility interface `t_sysconf(3NSL)` gets configurable XTI variables. The `t_sndreldata(3NSL)` interface initiates and responds to an orderly release with user data. The `t_rcvreldata(3NSL)` receives an orderly release indication or confirmation containing user data.

Note – The additional interfaces `t_sndreldata(3NSL)` and `t_rcvreldata(3NSL)` are used only with a specific transport called minimal OSI, which is not available on the Oracle Solaris platform. These interfaces are not available for use in conjunction with Internet Transports (TCP or UDP).

Packet Filtering Hooks

The packet filtering hooks interfaces help develop value added network solutions at the kernel level such as security (packet filtering and firewall) solutions and network address translation (NAT) solutions.

The packet filtering hooks interfaces provide the following capabilities:

- Notification each time a packet appears at one of the hook points
- Notification each time a new instance of IP is created to support a new zone booting that requires an exclusive instance of IP
- Kernel access to other basic network interface information such as interface names and addresses
- Interception of packets on the loopback interface

Loopback packet interception also provides access to packets as they move between zones that are using a shared instance of IP. This is the default model.

Packet Filtering Hooks Interfaces

Packet filtering hooks interfaces include kernel functions and data type definitions.

Packet Filtering Hooks Kernel Functions

The packet filtering hooks kernel functions are exported from the `misc/neti` and `misc/hook` kernel modules to support packet filtering. To use these functions, link your kernel modules with `-Nmisc/neti` and `-Nmisc/hook` so that the functions will be correctly loaded by the kernel.

`hook_alloc(9F)` Allocate a `hook_t` data structure.

`hook_free(9F)` Free a `hook_t` structure that was originally allocated by `hook_alloc()`.

<code>net_event_notify_register(9F)</code>	Register a function to be called when there is a change to a specified event.
<code>net_event_notify_unregister(9F)</code>	Indicate that there is no longer any desire to receive notification of changes to the specified event through calls to the specified callback function.
<code>net_getifname(9F)</code>	Retrieve the name given to the specified network interface.
<code>net_getlifaddr(9F)</code>	Retrieve the network address information for each specified logical interface.
<code>net_getmtu(9F)</code>	Retrieve information about the current MTU of the specified network interface.
<code>net_getpmtuenabled(9F)</code>	Indicate whether path MTU (PMTU) discovery is enabled for the specified network protocol.
<code>net_hook_register(9F)</code>	Add a hook that allows callbacks to be registered with events that belong to the specified network protocol.
<code>net_hook_unregister(9F)</code>	Disable callback hooks that were registered with <code>net_hook_register()</code> .
<code>net_inject(9F)</code>	Deliver network layer packets either into the kernel or onto the network.
<code>net_inject_alloc(9F)</code>	Allocate a <code>net_inject_t</code> structure.
<code>net_inject_free(9F)</code>	Free a <code>net_inject_t</code> structure that was originally allocated by <code>net_inject_alloc()</code> .
<code>net_instance_alloc(9F)</code>	Allocate a <code>net_instance_t</code> structure.
<code>net_instance_free(9F)</code>	Free a <code>net_instance_t</code> structure that was originally allocated by <code>net_instance_alloc()</code> .
<code>net_instance_notify_register(9F)</code>	Register the specified function to be called when there is a new instance added to or removed from the specified network instance.
<code>net_instance_notify_unregister(9F)</code>	Indicate that there is no longer any desire to receive notification of changes to the specified instance through calls to the specified callback function.
<code>net_instance_register(9F)</code>	Record the set of functions to be called when an event related to IP instance maintenance occurs.

<code>net_instance_unregister(9F)</code>	Remove the set of instances that were previously registered with <code>net_instance_register()</code> .
<code>net_ispartialchecksum(9F)</code>	Indicates whether the specified packet contains headers with only partial checksum values.
<code>net_isvalidchecksum(9F)</code>	Verify the layer 3 checksum and, in some cases, the layer 4 checksum in the specified packet.
<code>net_kstat_create(9F)</code>	Allocate and initialize a new <code>kstat(9S)</code> structure for the specified instance of IP.
<code>net_kstat_delete(9F)</code>	Remove a <code>kstat</code> for the specified instance of IP from the system.
<code>net_lifgetnext(9F)</code>	Search all of the logical interfaces that are associated with a physical network interface.
<code>net_phygetnext(9F)</code>	Search all of the network interfaces that a network protocol “owns.”
<code>net_phylookup(9F)</code>	Attempt to retrieve the specified interface name for a network protocol.
<code>net_protocol_lookup(9F)</code>	Locate an implementation of a network layer protocol.
<code>net_protocol_notify_register(9F)</code>	Register the specified function to be called when there is a change to the specified protocol.
<code>net_protocol_notify_unregister(9F)</code>	Remove the specified function from the list of functions to call.
<code>net_protocol_release(9F)</code>	Indicate that a reference to the specified network protocol is no longer required.
<code>net_routeto(9F)</code>	Indicate which network interface packets are sent.

Packet Filtering Hooks Data Types

The following types support the functions described above.

<code>hook_t(9S)</code>	A callback to be inserted into a networking event.
<code>hook_nic_event(9S)</code>	An event that has occurred and belongs to a network interface.
<code>hook_pkt_event(9S)</code>	A packet event structure passed through to hooks.
<code>net_inject_t(9S)</code>	Information about how to transmit a packet.

`net_instance_t(9S)` A collection of instances to be called when relevant events happen within IP.

Using the Packet Filtering Hooks Interfaces

A substantial amount of programming is required to work with the packet filtering hooks interfaces because this API supports multiple instances of the IP stack running concurrently in the same kernel. The IP stack allows multiple instances of itself for zones and multiple instances of the framework support packet interception in IP.

This section demonstrates the set up code to use the packet filtering hooks API to receive inbound IPv4 packets.

IP Instances

The first decision you need to make when you use this API is whether to accommodate multiple instances of IP running in the kernel or to only interact with the global zone.

To be aware of the presence of IP instances, register callback functions that are activated when an instance is created, destroyed, and shut down. Use `net_instance_alloc()` to allocate a `net_instance_t` packet event structure to store these three function pointers. Use `net_instance_free()` to free resources when you no longer need the callbacks and the structure. Specify `nin_name` to give the structure instance a name. Specify at least the `nin_create()` and `nin_destroy()` callbacks. The `nin_create()` function is called when a new instance of IP is created, and the `nin_destroy()` function is called when an instance of IP is destroyed.

Specifying `nin_shutdown()` is optional unless the code will be exporting information to `kstats`. To use `kstats` on a per-instance basis, use `net_kstat_create()` during the create callback. Cleanup of the `kstat` information must happen during the shutdown callback, not the destroy callback. Use `net_kstat_delete()` to clean up `kstat` information.

```
extern void *mycreate(const netid_t);

net_instance_t *n;

n = net_instance_alloc(NETINFO_VERSION);
if (n != NULL) {
    n->nin_create = mycreate;
    n->nin_destroy = mydestroy;
    n->nin_name = "my module";
    if (net_instance_register(n) != 0)
        net_instance_free(n);
}
```

If one or more instances of IP are present when `net_instance_alloc()` is called, the create callback will be called for each currently active instance. The framework that supports the callbacks ensures that only one of the create, destroy, or shutdown functions is active at any one time for a given instance. The framework also ensures that once the create callback has been called, the shutdown callback will only be called after create has completed. Similarly, the destroy callback does not start until the shutdown callback is complete.

The `mycreate()` function in the following example is a simple example of a create callback. The `mycreate()` function records the network instance identifier in its own private context structure and registers a new callback to be called when a new protocol (such as IPv4 or IPv6) is registered with this framework.

If no zones are running (and therefore no instances other than the global zone), calling `net_instance_register()` runs the create callback for the global zone. You must supply the destroy callback so that `net_instance_unregister()` can be called later. Attempts to call `net_instance_register()` with either the `nin_create` or `nin_destroy` fields set to NULL will fail.

```
void *
mycreate(const netid_t id)
{
    mytype_t *ctx;

    ctx = kmem_alloc(sizeof(*ctx), KM_SLEEP);
    ctx->instance_id = id;
    net_instance_notify_register(id, mynewproto, ctx);
    return (ctx);
}
```

The function `mynewproto()` should expect to be called each time a network protocol is either added to or removed from a networking instance. If registered network protocols are already operating within the given instance, then the create callback will be called for each protocol that already exists.

Protocol Registration

For this callback, only the `proto` argument is filled in by the caller. Neither an event nor a hook name can be meaningfully supplied at this point. In this example function, only events that announce the registration of the IPv4 protocol are being looked for.

The next step in this function is to discover when events are added to the IPv4 protocol by using the `net_protocol_notify_register()` interface to register the `mynewevent()` function.

```
static int
mynewproto(hook_notify_cmd_t cmd, void *arg, const char *proto,
           const char *event, const char *hook)
{
```

```

mytype_t *ctx = arg;

if (strcmp(proto, NHF_INET) != 0)
    return (0);

switch (cmd) {
    case HN_REGISTER :
        ctx->inet = net_protocol_lookup(s->id, proto);
        net_protocol_notify_register(s->inet, mynewevent, ctx);
        break;
    case HN_UNREGISTER :
    case HN_NONE :
        break;
}
return (0);
}

```

The table below lists all three protocols that could be expected to be seen with the `mynewproto()` callback. New protocols could be added in the future, so you must safely fail (return the value 0) any unknown protocols.

Programming Symbol	Protocol
NHF_INET	IPv4
NHF_INET6	IPv6
NHF_ARP	ARP

Event Registration

Just as the handling of instances and protocols is dynamic, the handling of the events that live under each protocol also is dynamic. Two types of events are supported by this API: network interface events and packet events.

In the function below, the announcement for the presence of the event for inbound packets for IPv4 is being checked for. When that announcement is seen, a `hook_t` structure is allocated, describing the function to be called for each inbound IPv4 packet.

```

static int
mynewevent(hook_notify_cmd_t cmd, void *arg, const char *parent,
           const char *event, const char *hook)
{
    mytype_t *ctx = arg;
    char buffer[32];
    hook_t *h;

    if ((strcmp(event, NH_PHYSICAL_IN) == 0) &&
        (strcmp(parent, NHF_INET) == 0)) {
        sprintf(buffer, sizeof(buffer), "mypkthook_%s_%s", parent, event);
        h = hook_alloc(HOOK_VERSION);
    }
}

```

```

        h->h_hint = HH_NONE;
        h->h_arg = s;
        h->h_name = strdup(buffer);
        h->h_func = mypkthook;
        s->hook_in = h;
        net_hook_register(ctx->inet, (char *)event, h);
    } else {
        h = NULL;
    }
    return (0);
}

```

The function `mynewevent()` will be called for each event that is added and removed. The following events are available.

Event Name	Data Structure	Comment
NH_PHYSICAL_IN	hook_pkt_event_t	This event is generated for every packet that arrives at the network protocol and has been received from a network interface driver.
NH_PHYSICAL_OUT	hook_pkt_event_t	This event is generated for every packet prior to delivery to the network interface driver for sending from the network protocol layer.
NH_FORWARDING	hook_pkt_event_t	This event is for all packets that have been received by the system and will be sent out another network interface. This event happens after NH_PHYSICAL_IN and before NH_PHYSICAL_OUT.
NH_LOOPBACK_IN	hook_pkt_event_t	This event is generated for packets that are received on the loopback interface or that are received by a zone that is sharing its network instance with the global zone.
NH_LOOPBACK_OUT	hook_pkt_event_t	This event is generated for packets that are sent on the loopback interface or that are being sent by a zone that is sharing its network instance with the global zone.
NH_NIC_EVENTS	hook_nic_event_t	This event is generated for specific changes of state for network interfaces.

For packet events, there is one specific event for each particular point in the IP stack. This is to enable you to be selective about exactly where in the flow of the packets you wish to intercept packets, without being overburdened by examining every packet event that happens inside the kernel. For network interface events the model is different, in part because the events are much lower in volume and because it is more likely that the developer will be interested in several of them, not just one.

The network interface event announces one of the following events:

- An interface is created (NE_PLUMB) or destroyed (NE_UNPLUMB).
- An interface changes state to up (NE_UP) or down (NE_DOWN).
- An interface has an address change (NE_ADDRESS_CHANGE).

New network interface events could be added in the future, so you must always return 0 for any unknown or unrecognized event that the callback function receives.

The Packet Hook

The packet hook function is called when a packet is received. In this case the function `mypkthook()` should expect to be called for each inbound packet that arrives in the kernel from a physical network interface. Packets generated internally, that flow between zones using the shared IP instance model or over the loopback interface, will not be seen.

To illustrate the difference between accepting a packet and allowing the function to return normally with what is required to drop a packet, the code below prints out the source and destination address of every 100th packet and then drops the packet, introducing a packet loss of 1%.

```
static int
mypkthook(hook_event_token_t tok, hook_data_t data, void *arg)
{
    static int counter = 0;
    mytupe_t *ctx = arg;
    hook_pkt_event_t *pkt = (hook_pkt_event_t)data;
    struct ip *ip;
    size_t bytes;

    bytes = msgdsize(pkt->hpe_mb);

    ip = (struct ip *)pkt->hpe_hdr;

    counter++;
    if (counter == 100) {
        printf("drop %d bytes received from %x to %x\n", bytes,
              ntohl(ip->ip_src.s_addr), ntohl(ip->ip_dst.s_addr));
        counter = 0;
        freemsg(*pkt->hpe_mp);
        *pkt->hpe_mp = NULL;
        pkt->hpe_mb = NULL;
        pkt->hpe_hdr = NULL;
        return (1);
    }
    return (0);
}
```

Packets received by this function, and all others that are called as a callback from a packet event, are received one at a time. There is no chaining together of packets with this interface, so you

should expect only one packet per call and expect `b_next` to always be `NULL`. While there is no other packet, a single packet may be comprised of several `mbk_t` structures chained together with `b_cont`.

Packet Filtering Hooks Example

Following is a complete example that can be compiled and loaded into the kernel.

Use the following commands to compile this code into a working kernel module on a 64-bit system:

```
# gcc -D_KERNEL -m64 -c full.c
# ld -dy -Nmisc/neti -Nmisc/hook -r full.o -o full
```

EXAMPLE 9-1 Packet Filtering Hooks Example Program

```
* Copyright (c) 2012, Oracle and/or its affiliates. All rights reserved.
*/
/*
* This file is a test module written to test the netinfo APIs in Oracle Solaris 11.
* It is being published to demonstrate how the APIs can be used.
*/
#include <sys/param.h>
#include <sys/sunddi.h>
#include <sys/modctl.h>
#include <sys/ddi.h>
#include "neti.h"

/*
* Module linkage information for the kernel.
*/
static struct modldrv modldrv = {
    &mod_miscops,      /* drv_modops */
    "neti test module", /* drv_linkinfo */
};

static struct modlinkage modlinkage = {
    MODREV_1,          /* ml_rev */
    &modldrv,          /* ml_linkage */
    NULL
};

typedef struct scratch_s {
    int          sentinel_1;
    netid_t      id;
    int          sentinel_2;
    int          event_notify;
    int          sentinel_3;
    int          v4_event_notify;
    int          sentinel_4;
    int          v6_event_notify;
    int          sentinel_5;
    int          arp_event_notify;
    int          sentinel_6;
};
```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```

int         v4_hook_notify;
int         sentinel_7;
int         v6_hook_notify;
int         sentinel_8;
int         arp_hook_notify;
int         sentinel_9;
hook_t     *v4_h_in;
int         sentinel_10;
hook_t     *v6_h_in;
int         sentinel_11;
hook_t     *arp_h_in;
int         sentinel_12;
net_handle_t v4;
int         sentinel_13;
net_handle_t v6;
int         sentinel_14;
net_handle_t arp;
int         sentinel_15;
} scratch_t;

#define MAX_RECALL_DOLOG      10000
char recall_myname[10];
net_instance_t *recall_global;
int recall_inited = 0;
int recall_doing[MAX_RECALL_DOLOG];
int recall_doidx = 0;
kmutex_t recall_lock;
int recall_continue = 1;
timeout_id_t recall_timeout;
int recall_steps = 0;
int recall_allocated = 0;
void *recall_alloclog[MAX_RECALL_DOLOG];
int recall_freed = 0;
void *recall_freelog[MAX_RECALL_DOLOG];

static int recall_init(void);
static void recall_fini(void);
static void *recall_create(const netid_t id);
static void recall_shutdown(const netid_t id, void *arg);
static void recall_destroy(const netid_t id, void *arg);
static int recall_newproto(hook_notify_cmd_t cmd, void *arg,
    const char *parent, const char *event, const char *hook);
static int recall_newevent(hook_notify_cmd_t cmd, void *arg,
    const char *parent, const char *event, const char *hook);
static int recall_newhook(hook_notify_cmd_t cmd, void *arg,
    const char *parent, const char *event, const char *hook);
static void recall_expire(void *arg);

static void recall_strfree(char *);
static char *recall_strdup(char *, int);

static void
recall_add_do(int mydo)
{
    mutex_enter(&recall_lock);
    recall_doing[recall_doidx] = mydo;

```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```

    recall_doidx++;
    recall_steps++;
    if ((recall_steps % 1000000) == 0)
        printf("stamp %d %d\n", recall_steps, recall_doidx);
    if (recall_doidx == MAX_RECALL_DOLOG)
        recall_doidx = 0;
    mutex_exit(&recall_lock);
}

static void *recall_alloc(size_t len, int wait)
{
    int i;

    mutex_enter(&recall_lock);
    i = recall_allocated++;
    if (recall_allocated == MAX_RECALL_DOLOG)
        recall_allocated = 0;
    mutex_exit(&recall_lock);

    recall_alloclog[i] = kmem_alloc(len, wait);
    return recall_alloclog[i];
}

static void recall_free(void *ptr, size_t len)
{
    int i;

    mutex_enter(&recall_lock);
    i = recall_freed++;
    if (recall_freed == MAX_RECALL_DOLOG)
        recall_freed = 0;
    mutex_exit(&recall_lock);

    recall_freelog[i] = ptr;
    kmem_free(ptr, len);
}

static void recall_assert(scratch_t *s)
{
    ASSERT(s->sentinel_1 == 0);
    ASSERT(s->sentinel_2 == 0);
    ASSERT(s->sentinel_3 == 0);
    ASSERT(s->sentinel_4 == 0);
    ASSERT(s->sentinel_5 == 0);
    ASSERT(s->sentinel_6 == 0);
    ASSERT(s->sentinel_7 == 0);
    ASSERT(s->sentinel_8 == 0);
    ASSERT(s->sentinel_9 == 0);
    ASSERT(s->sentinel_10 == 0);
    ASSERT(s->sentinel_11 == 0);
    ASSERT(s->sentinel_12 == 0);
    ASSERT(s->sentinel_13 == 0);
    ASSERT(s->sentinel_14 == 0);
    ASSERT(s->sentinel_15 == 0);
}

```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```
int
_init(void)
{
    int error;

    bzero(recall_doing, sizeof(recall_doing));
    mutex_init(&recall_lock, NULL, MUTEX_DRIVER, NULL);

    error = recall_init();
    if (error == DDI_SUCCESS) {
        error = mod_install(&modlinkage);
        if (error != 0)
            recall_fini();
    }

    recall_timeout = timeout(recall_expire, NULL, drv_usectohz(500000));

    return (error);
}

int
_fini(void)
{
    int error;

    recall_continue = 0;
    if (recall_timeout != NULL) {
        untimeout(recall_timeout);
        recall_timeout = NULL;
    }
    error = mod_remove(&modlinkage);
    if (error == 0) {
        recall_fini();
        delay(drv_usectohz(500000)); /* .5 seconds */

        mutex_destroy(&recall_lock);

        ASSERT(recall_inited == 0);
    }

    return (error);
}

int
_info(struct modinfo *info)
{
    return(0);
}

static int
recall_init()
{
    recall_global = net_instance_alloc(NETINFO_VERSION);

    strcpy(recall_myname, "full_");
    bcopy(((char *)&recall_global) + 4, recall_myname + 5, 4);
}
```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```

recall_myname[5] = (recall_myname[5] & 0x7f) | 0x20;
recall_myname[6] = (recall_myname[6] & 0x7f) | 0x20;
recall_myname[7] = (recall_myname[7] & 0x7f) | 0x20;
recall_myname[8] = (recall_myname[8] & 0x7f) | 0x20;
recall_myname[9] = '\0';

recall_global->nin_create = recall_create;
recall_global->nin_shutdown = recall_shutdown;
recall_global->nin_destroy = recall_destroy;
recall_global->nin_name = recall_myname;

if (net_instance_register(recall_global) != 0)
    return (DDI_FAILURE);

return (DDI_SUCCESS);
}

static void
recall_fini()
{
    if (recall_global != NULL) {
        net_instance_unregister(recall_global);
        net_instance_free(recall_global);
        recall_global = NULL;
    }
}

static void
recall_expire(void *arg)
{
    if (!recall_continue)
        return;

    recall_fini();

    if (!recall_continue)
        return;

    delay(drv_usecstohz(5000));    /* .005 seconds */

    if (!recall_continue)
        return;

    if (recall_init() == DDI_SUCCESS)
        recall_timeout = timeout(recall_expire, NULL,
            drv_usecstohz(5000));    /* .005 seconds */
}

static void *
recall_create(const netid_t id)
{
    scratch_t *s = kmem_zalloc(sizeof(*s), KM_SLEEP);

    if (s == NULL)

```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```
        return (NULL);

    recall_inited++;

    s->id = id;

    net_instance_notify_register(id, recall_newproto, s);

    return s;
}

static void
recall_shutdown(const netid_t id, void *arg)
{
    scratch_t *s = arg;

    ASSERT(s != NULL);
    recall_add_do(__LINE__);
    net_instance_notify_unregister(id, recall_newproto);

    if (s->v4 != NULL) {
        if (s->v4_h_in != NULL) {
            net_hook_unregister(s->v4, NH_PHYSICAL_IN,
                               s->v4_h_in);
            recall_strfree(s->v4_h_in->h_name);
            hook_free(s->v4_h_in);
            s->v4_h_in = NULL;
        }
        if (net_protocol_notify_unregister(s->v4, recall_newevent))
            cmn_err(CE_WARN,
                  "v4:net_protocol_notify_unregister(%p) failed",
                  s->v4);
        net_protocol_release(s->v4);
        s->v4 = NULL;
    }

    if (s->v6 != NULL) {
        if (s->v6_h_in != NULL) {
            net_hook_unregister(s->v6, NH_PHYSICAL_IN,
                               s->v6_h_in);
            recall_strfree(s->v6_h_in->h_name);
            hook_free(s->v6_h_in);
            s->v6_h_in = NULL;
        }
        if (net_protocol_notify_unregister(s->v6, recall_newevent))
            cmn_err(CE_WARN,
                  "v6:net_protocol_notify_unregister(%p) failed",
                  s->v6);
        net_protocol_release(s->v6);
        s->v6 = NULL;
    }

    if (s->arp != NULL) {
        if (s->arp_h_in != NULL) {
            net_hook_unregister(s->arp, NH_PHYSICAL_IN,
                               s->arp_h_in);
        }
    }
}
```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```

        recall_strfree(s->arp_h_in->h_name);
        hook_free(s->arp_h_in);
        s->arp_h_in = NULL;
    }
    if (net_protocol_notify_unregister(s->arp, recall_newevent))
        cmn_err(CE_WARN,
            "arp:net_protocol_notify_unregister(%p) failed",
            s->arp);
    net_protocol_release(s->arp);
    s->arp = NULL;
}

static void
recall_destroy(const netid_t id, void *arg)
{
    scratch_t *s = arg;

    ASSERT(s != NULL);

    recall_assert(s);

    ASSERT(s->v4 == NULL);
    ASSERT(s->v6 == NULL);
    ASSERT(s->arp == NULL);
    ASSERT(s->v4_h_in == NULL);
    ASSERT(s->v6_h_in == NULL);
    ASSERT(s->arp_h_in == NULL);
    kmem_free(s, sizeof(*s));

    ASSERT(recall_initd > 0);
    recall_initd--;
}

static int
recall_newproto(hook_notify_cmd_t cmd, void *arg, const char *parent,
    const char *event, const char *hook)
{
    scratch_t *s = arg;

    s->event_notify++;

    recall_assert(s);

    switch (cmd) {
    case HN_REGISTER :
        if (strcmp(parent, NHF_INET) == 0) {
            s->v4 = net_protocol_lookup(s->id, parent);
            net_protocol_notify_register(s->v4, recall_newevent, s);
        } else if (strcmp(parent, NHF_INET6) == 0) {
            s->v6 = net_protocol_lookup(s->id, parent);
            net_protocol_notify_register(s->v6, recall_newevent, s);
        } else if (strcmp(parent, NHF_ARP) == 0) {
            s->arp = net_protocol_lookup(s->id, parent);
            net_protocol_notify_register(s->arp, recall_newevent, s);
        }
    }
}

```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```
        break;

    case HN_UNREGISTER :
    case HN_NONE :
        break;
    }

    return 0;
}

static int
recall_do_event(hook_event_token_t tok, hook_data_t data, void *ctx)
{
    scratch_t *s = ctx;

    recall_assert(s);

    return (0);
}

static int
recall_newevent(hook_notify_cmd_t cmd, void *arg, const char *parent,
               const char *event, const char *hook)
{
    scratch_t *s = arg;
    char buffer[32];
    hook_t *h;

    recall_assert(s);

    if (strcmp(event, NH_PHYSICAL_IN) == 0) {

        sprintf(buffer, sizeof(buffer), "%s_%s_%s", recall_myname, parent, event);
        h = hook_alloc(HOOK_VERSION);
        h->h_hint = HH_NONE;
        h->h_arg = s;
        h->h_name = recall_strdup(buffer, KM_SLEEP);
        h->h_func = recall_do_event;
    } else {
        h = NULL;
    }

    if (strcmp(parent, NHF_INET) == 0) {
        s->v4_event_notify++;
        if (h != NULL) {
            s->v4_h_in = h;
            net_hook_register(s->v4, (char *)event, h);
        }
        net_event_notify_register(s->v4, (char *)event,
                                recall_newhook, s);
    } else if (strcmp(parent, NHF_INET6) == 0) {
        s->v6_event_notify++;
        if (h != NULL) {
            s->v6_h_in = h;
            net_hook_register(s->v6, (char *)event, h);
        }
    }
}
```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```

    }
    net_event_notify_register(s->v6, (char *)event,
        recall_newhook, s);

} else if (strcmp(parent, NHF_ARP) == 0) {
    s->arp_event_notify++;
    if (h != NULL) {
        s->arp_h_in = h;
        net_hook_register(s->arp, (char *)event, h);
    }
    net_event_notify_register(s->arp, (char *)event,
        recall_newhook, s);
}
recall_assert(s);

return (0);
}

static int
recall_newhook(hook_notify_cmd_t cmd, void *arg, const char *parent,
    const char *event, const char *hook)
{
    scratch_t *s = arg;

    recall_assert(s);

    if (strcmp(parent, NHF_INET) == 0) {
        s->v4_hook_notify++;
    } else if (strcmp(parent, NHF_INET6) == 0) {
        s->v6_hook_notify++;
    } else if (strcmp(parent, NHF_ARP) == 0) {
        s->arp_hook_notify++;
    }
    recall_assert(s);

    return (0);
}

static void recall_strfree(char *str)
{
    int len;

    if (str != NULL) {
        len = strlen(str);
        recall_free(str, len + 1);
    }
}

static char* recall_strdup(char *str, int wait)
{
    char *newstr;
    int len;

    len = strlen(str);

```

EXAMPLE 9-1 Packet Filtering Hooks Example Program (Continued)

```

newstr = recall_alloc(len, wait);
if (newstr != NULL)
    strcpy(newstr, str);

return (newstr);
}

```

EXAMPLE 9-2 net_inject Example Program

```

* Copyright (c) 2012, Oracle and/or its affiliates.
* All rights reserved.
*/

* PAMP driver - Ping Amplifier enables Solaris to send two ICMP echo
* responses for every ICMP request.
* This example provides a test module of the Oracle Solaris PF-hooks
* (netinfo(9f)) API. This example discovers ICMP echo
* implementation by intercepting inbound packets using
* physical-in' event hook.
* If the intercepted packet happens to be a ICMPv4 echo request,
* the module will generate a corresponding ICMP echo response
* which will then be sent to the network interface card using
* the net_inject(9f) function. The original ICMPv4 echo request will be
* allowed to enter the the IP stack so that the request can be
* processed by the destination IP stack.
* The destination stack in turn will send its own ICMPv4 echo response.
* Therefore there will be two ICMPv4 echo responses for a single
* ICMPv4 echo request.

*
* The following example code demonstrates two key functions of netinfo(9f) API:
*
* Packet Interception
*
* Packet Injection
*
* In order to be able to talk to netinfo(9f), the driver must allocate and
* register its own net_instance_t - 'pamp_ninst'. This happens in the
* pamp_attach() function, which implements 'ddi_attach' driver operation. The
* net_instance_t registers three callbacks with netinfo(9f) module:
*
*   _create
*   _shutdown
*   _destroy
*
* The netinfo(9f) command uses these functions to request the driver to
* create, shutdown, or destroy the driver context bound to a particular IP instance.
* This will enable the driver to handle packets for every IP stack found in
* the Oracle Solaris kernel. For purposes of this example, the driver is always
* implicitly bound to every IP instance.
*/

/* Use the following makefile to build the driver::
/* Begin Makefile */
ALL = pamp_drv pamp_drv.conf

pamp_drv = pamp_drv.o

```

EXAMPLE 9-2 net_inject Example Program (Continued)

```

pamp_drv.conf: pamp_drv
echo 'name="pamp_drv" parent="pseudo" instance=0;' > pamp_drv.conf

pamp_drv: pamp_drv.o
ld -dy -r -Ndrv/ip -Nmisc/neti -Nmsic/hook -o pamp_drv pamp_drv.o
pamp_drv.o: pamp_drv.c
cc -m64 -xmodel=kernel -D_KERNEL -c -o $@ $<

install:
cp pamp_drv /usr/kernel/drv/'isainfo -k'/pamp_drv
cp pamp_drv.conf /usr/kernel/drv/pamp_drv.conf

uninstall:
rm -rf /usr/kernel/drv/'isainfo -k'/pamp_drv
rm -rf /usr/kernel/drv/pamp_drv.conf

clean:
rm -f pamp_drv.o pamp_drv pamp_drv.conf

*End Makefile */

*
* The Makefile shown above will build a pamp_drv driver binary
* and pamp_drv.conf file for driver configuration. If you are
* building on a test machine, use 'make install' to place
* driver and configuration files in the specified location.
* Otherwise copy the pamp_drv binary and the pamp_drv.conf
* files to your test machine manually.
*
* Run the following command to load the driver to kernel:
    add_drv pam_drv
* Run the following command to unload the driver to kernel:
    rem_drv pamp_drv
*
* To check if your driver is working you need to use a snoop
* and 'ping' which will be running
* on a remote host. Start snoop on your network interface:
    snoop -d netX icmp
* Run a ping on a remote host:
ping -ns <test.box>
* test.box refers to the system where the driver is installed.

*
* The snoop should show there are two ICMP echo replies for every ICMP echo
* request. The expected output should be similar to the snoop output shown below:
* 172.16.1.2 -> 172.16.1.100 ICMP Echo request (ID: 16652 Sequence number: 0)
* 172.16.1.100 -> 172.16.1.2 ICMP Echo reply (ID: 16652 Sequence number: 0)
* 172.16.1.100 -> 172.16.1.2 ICMP Echo reply (ID: 16652 Sequence number: 0)
* 172.16.1.2 -> 172.16.1.100 ICMP Echo request (ID: 16652 Sequence number: 1)
* 172.16.1.100 -> 172.16.1.2 ICMP Echo reply (ID: 16652 Sequence number: 1)
* 172.16.1.100 -> 172.16.1.2 ICMP Echo reply (ID: 16652 Sequence number: 1)

```

EXAMPLE 9-2 net_inject Example Program (Continued)

```

* 172.16.1.2 -> 172.16.1.100 ICMP Echo request (ID: 16652 Sequence number: 2)
* 172.16.1.100 -> 172.16.1.2 ICMP Echo reply (ID: 16652 Sequence number: 2)
* 172.16.1.100 -> 172.16.1.2 ICMP Echo reply (ID: 16652 Sequence number: 2)
*/
#include <sys/atomic.h>
#include <sys/ksynch.h>
#include <sys/ddi.h>
#include <sys/modctl.h>
#include <sys/random.h>
#include <sys/sunddi.h>
#include <sys/stream.h>
#include <sys/devops.h>
#include <sys/stat.h>
#include <sys/modctl.h>
#include <sys/neti.h>
#include <sys/hook.h>
#include <sys/hook_event.h>
#include <sys/synch.h>
#include <inet/ip.h>
#include <netinet/in_system.h>
#include <netinet/in.h>
#include <netinet/ip.h>
#include <netinet/ip_icmp.h>

/*
 * This is a context for the driver. The context is allocated by
 * pamp_nin_create() callback for every IP instance found in kernel.
 */
typedef struct pamp_ipstack
{
    hook_t *pamp_phyin;
    int pamp_hook_ok;
    net_handle_t pamp_ipv4;
} pamp_ipstack_t;
static kmutex_t pamp_stcksmx;
/*
 * The netinstance, which passes driver callbacks to netinfo module.
 */
static net_instance_t *pamp_ninst = NULL;
/*
 * Solaris kernel driver APIs.
 */
static int pamp_getinfo(dev_info_t *, ddi_info_cmd_t, void *, void **);
static int pamp_attach(dev_info_t *, ddi_attach_cmd_t);
static int pamp_detach(dev_info_t *, ddi_detach_cmd_t); static dev_info_t *pamp_dev_info = NULL;
/*
 * Driver does not support any device operations.
 */

extern struct cb_ops no_cb_ops;

static struct dev_ops pamp_ops = {
    DEVO_REV,
    0,
    pamp_getinfo,

```

EXAMPLE 9-2 net_inject Example Program (Continued)

```

    nulldev,
    nulldev,
    pamp_attach,
    pamp_detach,
    nodev,
    &no_cb_ops,
    NULL,
    NULL,
    ddi_quiesce_not_needed,      /* quiesce */
};

static struct modldrv    pamp_module = {
&mod_driverops,
    "ECHO_1",
    &pamp_ops
};
static struct modlinkage pamp_modlink = {
    MODREV_1,
    &pamp_module,
    NULL
};

/*
 * Netinfo stack instance create/destroy/shutdown routines.
 */
static void *pamp_nin_create(const netid_t);
static void pamp_nin_destroy(const netid_t, void *);
static void pamp_nin_shutdown(const netid_t, void *);

/*
 * Callback to process intercepted packets delivered by hook event
 */
static int pamp_pkt_in(hook_event_token_t, hook_data_t, void *);

/*
 * Kernel driver getinfo operation
 */
static int
pamp_getinfo(dev_info_t *dip, ddi_info_cmd_t cmd, void * arg, void **resultp)
{
    int    e;

    switch (cmd) {
        case DDI_INFO_DEVT2DEVINFO:
            *resultp = pamp_dev_info;
            e = DDI_SUCCESS;
            break;
        case DDI_INFO_DEVT2INSTANCE:
            *resultp = NULL;
            e = DDI_SUCCESS;
            break;
        default:
            e = DDI_FAILURE;
    }

    return (e);
}

```

EXAMPLE 9-2 net_inject Example Program (Continued)

```

}
/*
 * Kernel driver attach operation. The job of the driver is to create a net
 * instance for our driver and register it with netinfo(9f)
 */
static int pamp_attach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    int    rc;
#define    RETURN(_x_)
    do    {
        mutex_exit(&pamp_stcksmx);
        return (_x_);
    } while (0)

    /*
     * Fail for all commands except DDI_ATTACH.
     */
    if (cmd != DDI_ATTACH) {
        return (DDI_FAILURE);
    }
    mutex_enter(&pamp_stcksmx);
    /*
     * It is an error to apply attach operation on a driver which is already
     * attached.
     */
    if (pamp_ninst != NULL) {
        RETURN(DDI_FAILURE);
    }
    /*
     * At most one driver instance is allowed (instance 0).
     */
    if (ddi_get_instance(dip) != 0) {
        RETURN(DDI_FAILURE);
    }

    rc = ddi_create_minor_node(dip, "pamp", S_IFCHR, 0, DDI_PSEUDO, 0);
    if (rc != DDI_SUCCESS) {
        ddi_remove_minor_node(dip, NULL);
        RETURN(DDI_FAILURE);
    }

    /*
     * Create and register pamp net instance. Note we are assigning
     * callbacks _create, _destroy, _shutdown. These callbacks will ask
     * our driver to create/destroy/shutdown our IP driver instances.
     */
    pamp_ninst = net_instance_alloc(NETINFO_VERSION);
    if (pamp_ninst == NULL) {
        ddi_remove_minor_node(dip, NULL);
        RETURN(DDI_FAILURE);
    }

    pamp_ninst->nin_name = "pamp";
    pamp_ninst->nin_create = pamp_nin_create;
    pamp_ninst->nin_destroy = pamp_nin_destroy;
    pamp_ninst->nin_shutdown = pamp_nin_shutdown;

```

EXAMPLE 9-2 net_inject Example Program (Continued)

```

pamp_dev_info = dip;
mutex_exit(&pamp_stcksmx);

/*
 * Although it is not shown in the following example, it is
 * recommended that all mutexes/exclusive locks be released before *
 * calling net_instance_register(9F) to avoid a recursive lock
 * entry. As soon as pamp_ninst is registered, the
 * net_instance_register(9f) will call pamp_nin_create() callback.
 * The callback will run in the same context as the one in which
 * pamp_attach() is running. If pamp_nin_create() grabs the same
 * lock held already by pamp_attach(), then such a lock is being
 * operated on recursively.
 */
(void) net_instance_register(pamp_ninst);

return (DDI_SUCCESS);
#undef RETURN
}

/*
 * The detach function will unregister and destroy our driver netinstance. The same rules
 * for exclusive locks/mutexes introduced for attach operation apply to detach.
 * The netinfo will take care to call the shutdown()/destroy() callbacks for
 * every IP stack instance.
 */
static int
pamp_detach(dev_info_t *dip, ddi_detach_cmd_t cmd)
{
    pamp_ipstack_t    *pamp_ipstack;
    net_instance_t    *ninst = NULL;

    /*
     * It is an error to apply detach operation on driver, when another
     * detach operation is running (in progress), or when detach operation
     * is complete (pamp_ninst).
     */
    mutex_enter(&pamp_stcksmx);
    if (pamp_ninst == NULL) {
        mutex_exit(&pamp_stcksmx);
        return (DDI_FAILURE);
    }

    ninst = pamp_ninst;
    pamp_ninst = NULL;
    mutex_exit(&pamp_stcksmx);

    /*
     * Calling net_instance_unregister(9f) will invoke pamp_nin_destroy()
     * for every pamp_ipstack instance created so far. Therefore it is advisable
     * to not hold any mutexes, because it might get grabbed by pamp_nin_destroy() function.
     */
    net_instance_unregister(ninst);
    net_instance_free(ninst);

    (void) ddi_get_instance(dip);

```

EXAMPLE 9-2 net_inject Example Program (Continued)

```
    ddi_remove_minor_node(dip, NULL);

    return (DDI_SUCCESS);
}

/*
 * Netinfo callback, which is supposed to create an IP stack context for our
 * ICMP echo server.
 *
 * NOTE: NULL return value is not interpreted as a failure here. The
 * pamp_nin_shutdown()/pamp_nin_destroy() will receive NULL pointer for IP stack
 * instance with given 'netid' id.
 */
static void *
pamp_nin_create(const netid_t netid)
{
    pamp_ipstack_t    *pamp_ipstack;

    pamp_ipstack = (pamp_ipstack_t *)kmem_zalloc(
        sizeof (pamp_ipstack_t), KM_NOSLEEP);

    if (pamp_ipstack == NULL) {
        return (NULL);
    }

    HOOK_INIT(pamp_ipstack->pamp_phyin, pamp_pkt_in, "pkt_in",
        pamp_ipstack);

    pamp_ipstack->pamp_ipv4 = net_protocol_lookup(netid, NHF_INET);
    if (pamp_ipstack->pamp_ipv4 == NULL) {
        kmem_free(pamp_ipstack, sizeof (pamp_ipstack_t));
        return (NULL);
    }

    pamp_ipstack->pamp_hook_ok = net_hook_register(
        pamp_ipstack->pamp_ipv4, NH_PHYSICAL_IN, pamp_ipstack->pamp_phyin);
    if (pamp_ipstack->pamp_hook_ok != 0) {
        net_protocol_release(pamp_ipstack->pamp_ipv4);
        hook_free(pamp_ipstack->pamp_phyin);
        kmem_free(pamp_ipstack, sizeof (pamp_ipstack_t));
        return (NULL);
    }

    return (pamp_ipstack);
}

/*
 * This event is delivered right before the particular stack instance is
 * destroyed.
 */
static void
pamp_nin_shutdown(const netid_t netid, void *stack)
{
    return;
}
```

EXAMPLE 9-2 net_inject Example Program (Continued)

```

}

/*
 * Important to note here that the netinfo(9f) module ensures that no
 * no pamp_pkt_in() is "running" when the stack it is bound to is being destroyed.
 */

static void
pamp_nin_destroy(const netid_t netid, void *stack)
{
    pamp_ipstack_t    *pamp_ipstack = (pamp_ipstack_t *)stack;

    /*
     * Remember stack can be NULL! The pamp_nin_create() function returns
     * NULL on failure. The return value of pamp_nin_create() function will
     * be 'kept' in netinfo module as a driver context for particular IP
     * instance. As soon as the instance is destroyed the NULL value
     * will appear here in pamp_nin_destroy(). Same applies to
     * pamp_nin_shutdown(). Therefore our driver must be able to handle
     * NULL here.
     */
    if (pamp_ipstack == NULL)
        return;

    /*
     * If driver has managed to initialize packet hook, then it has to be
     * unhooked here.
     */
    if (pamp_ipstack->pamp_hook_ok != -1) {
        (void) net_hook_unregister(pamp_ipstack->pamp_ipv4,
            NH_PHYSICAL_IN, pamp_ipstack->pamp_phyin);
        hook_free(pamp_ipstack->pamp_phyin);
        (void) net_protocol_release(pamp_ipstack->pamp_ipv4);
    }

    kmem_free(pamp_ipstack, sizeof (pamp_ipstack_t));
}

/*
 * Packet hook handler
 *
 * Function receives intercepted IPv4 packets coming from NIC to IP stack. If
 * inbound packet is ICMP echo request, then function will generate ICMP echo
 * response and use net_inject() to send it to network. Function will also let
 * ICMP echo request in, so it will be still processed by destination IP stack,
 * which should also generate its own ICMP echo response. The snoop should show
 * you there will be two ICMP echo responses leaving the system where the pamp
 * driver is installed
 */

static int
pamp_pkt_in(hook_event_token_t ev, hook_data_t info, void *arg)
{
    hook_pkt_event_t    *hpe = (hook_pkt_event_t *)info;
    phy_if_t            phyif;

```

EXAMPLE 9-2 net_inject Example Program (Continued)

```

struct ip      *ip;

/*
 * Since our pamp_pkt_in callback is hooked to PHYSICAL_IN hook pkt.
 * event only, the physical interface index will always be passed as
 * hpe_ifp member.
 *
 * If our hook processes PHYSICAL_OUT hook pkt event, then
 * the physical interface index will be passed as hpe_ofp member.
 */
phyif = hpe->hpe_ifp;

ip = hpe->hpe_hdr;
if (ip->ip_p == IPPROTO_ICMP) {
    mblk_t      *mb;

    /*
     * All packets are copied/placed into a continuous buffer to make
     * parsing easier.
     */
    if ((mb = msgpullup(hpe->hpe_mb, -1)) != NULL) {
        struct icmp      *icmp;
        pamp_ipstack_t    *pamp_ipstack = (pamp_ipstack_t *)arg;

        ip = (struct ip *)mb->b_rptr;
        icmp = (struct icmp *)mb->b_rptr + IPH_HDR_LENGTH(ip);

        if (icmp->icmp_type == ICMP_ECHO) {
            struct in_addr  addr;
            uint32_t        sum;
            mblk_t          *echo_resp = copymsg(mb);
            net_inject_t     ninj;

            /*
             * We need to make copy of packet, since we are
             * going to turn it into ICMP echo response.
             */
            if (echo_resp == NULL) {
                return (0);
            }
            ip = (struct ip *)echo_resp->b_rptr;
            addr = ip->ip_src;
            ip->ip_src = ip->ip_dst;
            ip->ip_dst = addr;
            icmp = (struct icmp *) (echo_resp->b_rptr + IPH_HDR_LENGTH(ip));
            icmp->icmp_type = ICMP_ECHO_REPLY;
            sum = ntohs(icmp->icmp_cksum) & 0xffff;
            sum += (ICMP_ECHO_REQUEST - ICMP_ECHO_REPLY);
            icmp->icmp_cksum =
                htons(~((sum >> 16) + (sum & 0xffff)));

            /*
             * Now we have assembled an ICMP response with
             * correct chksum. It's time to send it out.
             * We have to initialize command for
             * net_inject(9f) -- ninj.
             */

```

EXAMPLE 9-2 net_inject Example Program (Continued)

```

*/
ninj.ni_packet = echo_resp;
ninj.ni_physical = phyif;
/*
 * As we are going use NI_QUEUE_OUT to send
 * our ICMP response, we don't need to set up
 * .ni_addr, which is required for NI_DIRECT_OUT
 * injection path only. In such case packet
 * bypasses IP stack routing and is pushed
 * directly to physical device queue. Therefore
 * net_inject(9f) requires as to specify
 * next-hop IP address.
 *
 * Using NI_QUEUE_OUT is more convenient for us
 * since IP stack will take care of routing
 * process and will find out 'ni_addr'
 * (next-hop) address on its own.
 */
(void) net_inject(pamp_ipstack->pamp_ipv4,
    NI_QUEUE_OUT, &ninj);
    }
}

/*
 * 0 as return value will let packet in.
 */
return (0);
}

/*
 * Kernel module handling.
 */
int init()
{
    mutex_init(&pamp_stcksmx, "pamp_mutex", MUTEX_DRIVER, NULL);
    return (mod_install(&pamp_modlink));
}

int fini()
{
    int rv;

    rv = mod_remove(&pamp_modlink);
    return (rv);
}

int info(struct modinfo *modinfop)
{
    return (mod_info(&pamp_modlink, modinfop));
}

```


Transport Selection and Name-to-Address Mapping

This chapter describes selecting transports and resolving network addresses. This chapter further describes interfaces that enable you to specify the available communication protocols for an application. The chapter also explains additional interfaces that provide direct mapping of names to network addresses.

- “Transport Selection” on page 243
- “Name-to-Address Mapping” on page 244

Note – In this chapter, the terms *network* and *transport* are used interchangeably. The terms refer to the programmatic interface that conforms to the transport layer of the OSI Reference Mode. The term *network* is also used to refer to the physical collection of computers that are connected through some electronic medium.

Transport Selection



Caution – The interfaces that are described in this chapter are multithread safe. “Multithread safe” means that you can use applications that contain transport selection interface calls freely in a multithreaded application. These interface calls do not provide linear scalability because the calls are not re-entrant.

A distributed application must use a standard interface to the transport services to be portable to different protocols. Transport selection services provide an interface that allows an application to select which protocols to use. This interface makes an application independent of protocol and medium.

Transport selection means that a client application can easily try each available transport until the client establishes communication with a server. Transport selection enables request acceptance on multiple transports by server applications. The applications can then

communicate over a number of protocols. Transports can be tried in either the order specified by the local default sequence or in an order specified by the user.

Choosing from the available transports is the responsibility of the application. The transport selection mechanism makes that selection uniform and simple.

Name-to-Address Mapping

Name-to-address mapping enables an application to obtain the address of a service on a specified host independent of the transport used. Name-to-address mapping consists of the following interfaces:

<code>netdir_getbyname(3NSL)</code>	Maps the host and service name to a set of addresses
<code>netdir_getbyaddr(3NSL)</code>	Maps addresses into host and service names
<code>netdir_free(3NSL)</code>	Frees structures allocated by the name-to-address translation routines
<code>taddr2uaddr(3NSL)</code>	Translates an address and returns a transport-independent character representation of the address
<code>uaddr2taddr(3NSL)</code>	The universal address is translated into a <code>netbuf</code> structure
<code>netdir_options(3NSL)</code>	Interfaces to transport-specific capabilities such as the broadcast address and reserved port facilities of TCP and UDP
<code>netdir_perror(3NSL)</code>	Displays a message stating why one of the routines that map name-to-address failed on <code>stderr</code> .
<code>netdir_spererror(3NSL)</code>	Returns a string containing the error message stating why one of the routines that map name-to-address failed.

The first argument of each routine points to a `netconfig(4)` structure that describes a transport. The routine uses the array of directory-lookup library paths in the `netconfig(4)` structure to call each path until the translation succeeds.

The name-to-address libraries are described in [Table 10–1](#). The routines that are described in “Using the Name-to-Address Mapping Routines” on [page 246](#) are defined in the `netdir(3NSL)` man page.

Note – The following libraries no longer exist in the Oracle Solaris environment: `tcpip.so`, `switch.so`, and `nis.so`. For more information on this change, see the `nsswitch.conf(4)` man page and the NOTES section of the `gethostbyname(3NSL)` man page.

TABLE 10-1 Name-to-Address Libraries

Library	Transport Family	Description
-	inet	The name-to-address mapping for networks of the protocol family <code>inet</code> is provided by the name service switch based on the entries for <code>hosts</code> and <code>services</code> in the file <code>nsswitch.conf(4)</code> . For networks of other families, the dash indicates a nonfunctional name-to-address mapping.

straddr.so Library

Name-to-address translation files for the `straddr.so` library are created by the system administrator. The system administrator also maintains these translation files. The `straddr.so` files are `/etc/net/transport-name/hosts` and `/etc/net/transport-name/services`. `transport-name` is the local name of the transport that accepts string addresses, which is specified in the `network ID` field of the `/etc/netconfig` file. For example, the host file for `ticlts` would be `/etc/net/ticlts/hosts`, and the service file for `ticlts` would be `/etc/net/ticlts/services`.

Most string addresses do not distinguish between *host* and *service*. However, separating the string into a host part and a service part is consistent with other transports. The `/etc/net/transport-name/hosts` file contains a text string that is assumed to be the host address, followed by the host name:

```
joyluckaddr    joyluck
carpediemaddr  carpediem
thehopaddr     thehop
pongoaddr      pongo
```

The `/etc/net/transport-name/services` file contains service names followed by strings that identify the service address:

```
rpcbind    rpc
listen     serve
```

The routines create the full-string address by concatenating the host address, a period (`.`), and the service address. For example, the address of the `listen` service on `pongo` is `pongoaddr.serve`.

When an application requests the address of a service on a particular host on a transport that uses this library, the host name must be in `/etc/net/transport/hosts`. The service name must be in `/etc/net/transport/services`. If either name is missing, the name-to-address translation fails.

Using the Name-to-Address Mapping Routines

This section is an overview of the mapping routines that are available for use. The routines return or convert the network names to their respective network addresses. Note that `netdir_getbyname(3NSL)`, `netdir_getbyaddr(3NSL)`, and `taddr2uaddr(3NSL)` return pointers to data that must be freed by calls to `netdir_free(3NSL)`.

```
int netdir_getbyname(struct netconfig *nconf,
                   struct nd_hostserv *service, struct nd_addrlist **addrs);
```

`netdir_getbyname(3NSL)` maps the host and service name specified in *service* to a set of addresses that are consistent with the transport identified in *nconf*. The `nd_hostserv` and `nd_addrlist` structures are defined in the `netdir(3NSL)` man page. A pointer to the addresses is returned in *addrs*.

To find all addresses of a host and service on all available transports, call `netdir_getbyname(3NSL)` with each `netconfig(4)` structure returned by either `getnetpath(3NSL)` or `getnetconfig(3NSL)`.

```
int netdir_getbyaddr(struct netconfig *nconf,
                   struct nd_hostservlist **service, struct netbuf *netaddr);
```

`netdir_getbyaddr(3NSL)` maps addresses into host and service names. The interface is called with an address in *netaddr* and returns a list of host-name and service-name pairs in *service*. The `nd_hostservlist` structure is defined in `netdir(3NSL)`.

```
void netdir_free(void *ptr, int struct_type);
```

The `netdir_free(3NSL)` routine frees structures allocated by the name-to-address translation routines. The parameters can take the values that are shown in the following table.

TABLE 10-2 `netdir_free(3NSL)` Routines

struct_type	ptr
ND_HOSTSERV	Pointer to an <code>nd_hostserv</code> structure
ND_HOSTSERVLIST	Pointer to an <code>nd_hostservlist</code> structure
ND_ADDR	Pointer to a <code>netbuf</code> structure
ND_ADDRLIST	Pointer to an <code>nd_addrlist</code> structure

```
char *taddr2uaddr(struct netconfig *nconf, struct netbuf *addr);
```

`taddr2uaddr(3NSL)` translates the address pointed to by *addr* and returns a transport-independent character representation of the address. This character representation is

called a universal address. The value that is given in *nconf* specifies the transport for which the address is valid. The universal address can be freed by `free(3C)`.

```
struct netbuf *uaddr2taddr(struct netconfig *nconf, char *uaddr);
```

The universal address pointed to by *uaddr* is translated into a `netbuf` structure. *nconf* specifies the transport for which the address is valid.

```
int netdir_options(const struct netconfig *config,
                  const int option, const int fildes, char *point_to_args);
```

`netdir_options(3NSL)` provides interfaces to transport-specific capabilities, such as the broadcast address and reserved port facilities of TCP and UDP. The value of *nconf* specifies a transport, while *option* specifies the transport-specific action to take. The value in *option* might disable consideration of the value in *fd*. The fourth argument points to operation-specific data.

The following table shows the values used for *option*.

TABLE 10-3 Values for `netdir_options`

Option	Description
ND_SET_BROADCAST	Sets the transport for broadcast if the transport supports broadcast
ND_SET_RESERVEDPORT	Enables application binding to reserved ports if allowed by the transport
ND_CHECK_RESERVEDPORT	Verifies that an address corresponds to a reserved port if the transport supports reserved ports
ND_MERGEADDR	Transforms a locally meaningful address into an address to which client hosts can connect

The `netdir_perror(3NSL)` routine displays a message stating why one of the routines that map name-to-address failed on `stderr`.

```
void netdir_perror(char *s);
```

The `netdir_sperror(3NSL)` routine returns a string containing the error message stating why one of the routines that map name-to-address failed.

```
char *netdir_sperror(void);
```

The following example shows network selection and name-to-address mapping.

EXAMPLE 10-1 Network Selection and Name-to-Address Mapping

```
#include <netconfig.h>
#include <netdir.h>
#include <sys/tiuser.h>
```

EXAMPLE 10-1 Network Selection and Name-to-Address Mapping (Continued)

```

struct nd_hostserv nd_hostserv; /* host and service information */
struct nd_addrlist *nd_addrlistp; /* addresses for the service */
struct netbuf *netbufp; /* the address of the service */
struct netconfig *nconf; /* transport information*/
int i; /* the number of addresses */
char *uaddr; /* service universal address */
void *handlep; /* a handle into network selection */
/*
 * Set the host structure to reference the "date"
 * service on host "gandalf"
 */
nd_hostserv.h_host = "gandalf";
nd_hostserv.h_serv = "date";
/*
 * Initialize the network selection mechanism.
 */
if ((handlep = setnetpath()) == (void *)NULL) {
    nc_perror(argv[0]);
    exit(1);
}
/*
 * Loop through the transport providers.
 */
while ((nconf = getnetpath(handlep)) != (struct netconfig *)NULL)
{
    /*
     * Print out the information associated with the
     * transport provider described in the "netconfig"
     * structure.
     */
    printf("Transport provider name: %s\n", nconf->nc_netid);
    printf("Transport protocol family: %s\n", nconf->nc_protofmly);
    printf("The transport device file: %s\n", nconf->nc_device);
    printf("Transport provider semantics: ");
    switch (nconf->nc_semantics) {
    case NC_TPI_COTS:
        printf("virtual circuit\n");
        break;
    case NC_TPI_COTS_ORD:
        printf("virtual circuit with orderly release\n");
        break;

    case NC_TPI_CLTS:
        printf("datagram\n");
        break;
    }
    /*
     * Get the address for service "date" on the host
     * named "gandalf" over the transport provider
     * specified in the netconfig structure.
     */
    if (netdir_getbyname(nconf, &nd_hostserv, &nd_addrlistp) != ND_OK) {
        printf("Cannot determine address for service\n");
        netdir_perror(argv[0]);
        continue;
    }
}

```

EXAMPLE 10-1 Network Selection and Name-to-Address Mapping (Continued)

```
}
printf("<%d> addresses of date service on gandalf:\n",
      nd_addrlistp->n_cnt);
/*
 * Print out all addresses for service "date" on
 * host "gandalf" on current transport provider.
 */
netbufp = nd_addrlistp->n_addrs;
for (i = 0; i < nd_addrlistp->n_cnt; i++, netbufp++) {
    uaddr = taddr2uaddr(nconf,netbufp);
    printf("%s\n",uaddr);
    free(uaddr);
}
netdir_free( nd_addrlistp, ND_ADDRLIST );
}
endnetconfig(handlep);
```


Real-time Programming and Administration

This chapter describes writing and porting real-time applications to run under SunOS. This chapter is written for programmers that are experienced in writing real-time applications and for administrators familiar with real-time processing and the Oracle Solaris system.

This chapter discusses the following topics:

- Scheduling needs of real-time applications, which are covered in [“The Real-Time Scheduler” on page 255](#).
- [“Memory Locking” on page 265](#).
- [“Asynchronous Network Communication” on page 272](#).

Basic Rules of Real-time Applications

Real-time response is guaranteed when certain conditions are met. This section identifies these conditions and some of the more significant design errors.

Most of the potential problems described here can degrade the response time of the system. One of the potential problems can freeze a workstation. Other, more subtle, mistakes are priority inversion and system overload.

An Oracle Solaris real-time process has the following characteristics:

- Runs in the RT scheduling class, as described in [“The Real-Time Scheduler” on page 255](#)
- Locks down all the memory in its process address space, as described in [“Memory Locking” on page 265](#)
- Is from a program in which all dynamic binding is completed early, as described in [“Shared Libraries” on page 253](#)

Real-time operations are described in this chapter in terms of single-threaded processes, but the description can also apply to multithreaded processes. For detailed information about multithreaded processes, see the [Multithreaded Programming Guide](#). To guarantee real-time

scheduling of a thread, the thread must be created as a bound thread. Furthermore, the thread's LWP must be run in the RT scheduling class. The locking of memory and early dynamic binding is effective for all threads in a process.

When a process is the highest priority real-time process, the process acquires the processor within the guaranteed dispatch latency period of becoming runnable. For more information, see [“Dispatch Latency” on page 255](#). The process continues to run for as long as it remains the highest priority runnable process.

A real-time process can lose control of the processor because of other system events. A real-time process can also be unable to gain control of the processor because of other system events. These events include external events, such as interrupts, resource starvation, waiting on external events such as synchronous I/O, and preemption by a higher priority process.

Real-time scheduling generally does not apply to system initialization and termination services such as `open(2)` and `close(2)`.

Factors that Degrade Response Time

The problems described in this section all increase the response time of the system to varying extents. The degradation can be serious enough to cause an application to miss a critical deadline.

Real-time processing can also impair the operation of aspects of other applications that are active on a system that is running a real-time application. Because real-time processes have higher priority, time-sharing processes can be prevented from running for significant amounts of time. This phenomenon can cause interactive activities, such as displays and keyboard response time, to slow noticeably.

Synchronous I/O Calls

System response under SunOS provides no bounds to the timing of I/O events. This means that synchronous I/O calls should never be included in any program segment whose execution is time critical. Even program segments that permit very large time bounds must not perform synchronous I/O. Mass storage I/O is such a case, where causing a read or write operation hangs the system while the operation takes place.

A common application mistake is to perform I/O to get error message text from disk. Performing I/O in this fashion should be done from an independent process or independent thread. This independent process or independent thread should not run in real time.

Interrupt Servicing

Interrupt priorities are independent of process priorities. The priorities that are set for a group of processes are not inherited by the services of hardware interrupts that result from those processes' actions. As a consequence, devices controlled by high-priority real-time processes do not necessarily have high-priority interrupt processing.

Shared Libraries

Time-sharing processes can save significant amounts of memory by using dynamically linked, shared libraries. This type of linking is implemented through a form of file mapping. Dynamically linked library routines cause implicit reads.

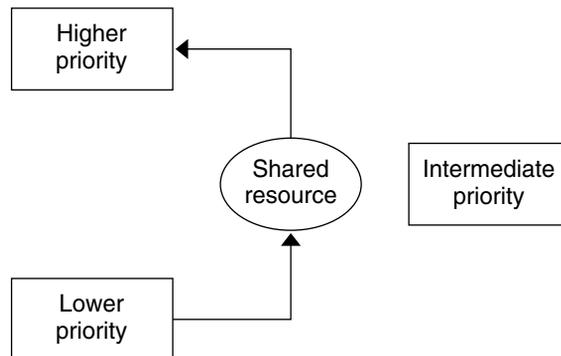
Real-time programs can set the environment variable `LD_BIND_NOW` to a non-NULL value when the program is invoked. Setting the value of this environment value allows the use of shared libraries while avoiding dynamic binding. This procedure also forces all dynamic linking to be bound before the program begins execution. See the *Linker and Libraries Guide* for more information.

Priority Inversion

A time-sharing process can block a real-time process by acquiring a resource that is required by a real-time process. Priority inversion occurs when a higher priority process is blocked by a lower priority process. The term *blocking* describes a situation in which a process must wait for one or more processes to relinquish control of resources. Real-time processes might miss their deadlines if this blocking is prolonged.

Consider the case that is depicted in the following figure, where a high-priority process requires a shared resource. A lower priority process holds the resource and is preempted by an intermediate priority process, blocking the high-priority process. Any number of intermediate processes can be involved. All intermediate processes must finish executing, as well as the lower-priority process' critical section. This series of executions can take an arbitrarily long time.

FIGURE 11-1 Unbounded Priority Inversion



This issue and the methods of dealing with this issue are described in “[Mutual Exclusion Lock Attributes](#)” in *Multithreaded Programming Guide*.

Sticky Locks

A page is permanently locked into memory when its lock count reaches 65535 (0xFFFF). The value 0xFFFF is defined by the implementation and might change in future releases. Pages that are locked this way cannot be unlocked.

Runaway Real-time Processes

Runaway real-time processes can cause the system to halt. Such runaway processes can also slow the system response so much that the system appears to halt.

Note – If you have a runaway process on a SPARC system, press Stop-A. You might have to do press Stop-A more than one time. If pressing Stop-A does not work, turn the power off, wait a moment, then turn the power back on. If you have a runaway process on a non-SPARC system, turn the power off, wait a moment, then turn the power back on.

When a high priority real-time process does not relinquish control of the CPU, you must break the infinite loop in order to regain control of the system. Such a runaway process does not respond to Control-C. Attempts to use a shell set at a higher priority than the priority of the runaway process do not work.

Asynchronous I/O Behavior

Asynchronous I/O operations do not always execute in the sequence in which the operations are queued to the kernel. Asynchronous operations do not necessarily return to the caller in the sequence in which the operations were performed.

If a single buffer is specified for a rapid sequence of calls to `aio_read`, the buffer's state is uncertain. The uncertainty of the buffer's state is from the time the first call is made to the time the last result is signaled to the caller.

An individual `aio_result_t` structure can be used for only one asynchronous operation. The operation can be a read or a write operation.

Real-time Files

SunOS provides no facilities to ensure that files are allocated as physically contiguous.

For regular files, the `read(2)` and `write(2)` operations are always buffered. An application can use `mmap(2)` and `msync(3C)` to effect direct I/O transfers between secondary storage and process memory.

The Real-Time Scheduler

Real-time scheduling constraints are necessary to manage data acquisition or process control hardware. The real-time environment requires that a process be able to react to external events in a bounded amount of time. Such constraints can exceed the capabilities of a kernel that is designed to provide a fair distribution of the processing resources to a set of time-sharing processes.

This section describes the SunOS real-time scheduler, its priority queue, and how to use system calls and utilities that control scheduling.

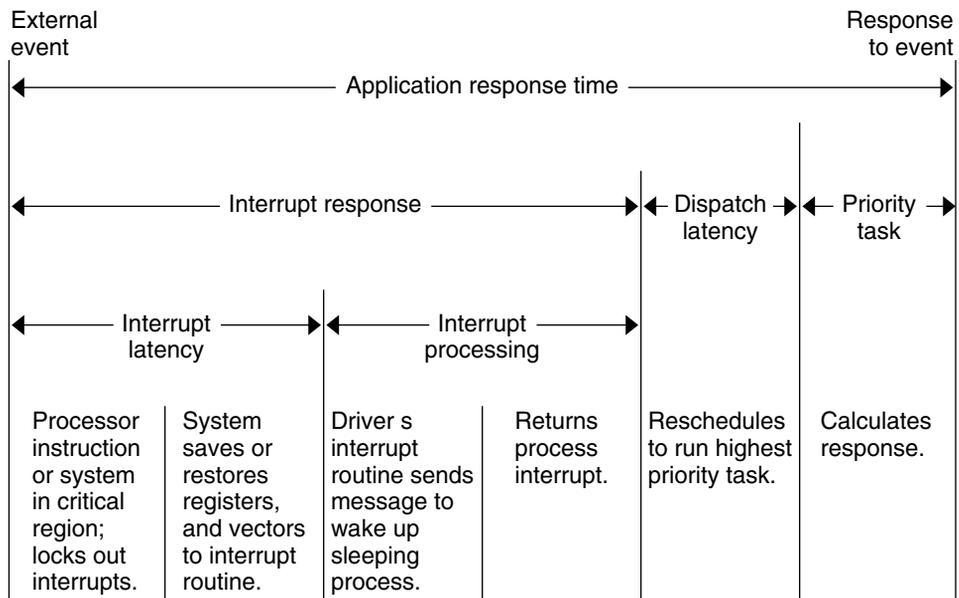
Dispatch Latency

The most significant element in scheduling behavior for real-time applications is the provision of a real-time scheduling class. The standard time-sharing scheduling class is not suitable for real-time applications because this scheduling class treats every process equally. The standard time-sharing scheduling class has a limited notion of priority. Real-time applications require a scheduling class in which process priorities are taken as absolute. Real-time applications also require a scheduling class in which process priorities are changed only by explicit application operations.

The term *dispatch latency* describes the amount of time a system takes to respond to a request for a process to begin operation. With a scheduler that is written specifically to honor application priorities, real-time applications can be developed with a bounded dispatch latency.

The following figure illustrates the amount of time an application takes to respond to a request from an external event.

FIGURE 11-2 Application Response Time



The overall application response time consists of the interrupt response time, the dispatch latency, and the application's response time.

The interrupt response time for an application includes both the interrupt latency of the system and the device driver's own interrupt processing time. The interrupt latency is determined by the longest interval that the system must run with interrupts disabled. This time is minimized in SunOS using synchronization primitives that do not commonly require a raised processor interrupt level.

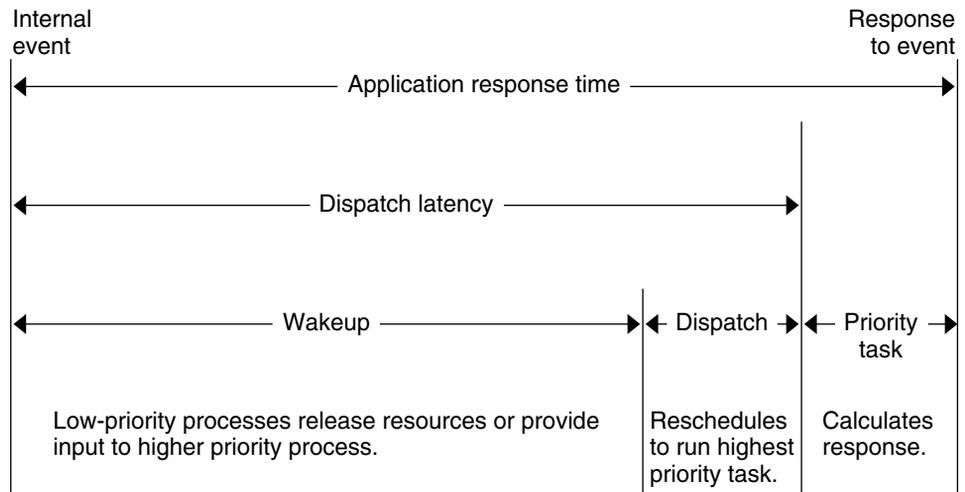
During interrupt processing, the driver's interrupt routine wakes the high-priority process and returns when finished. The system detects that a process with higher priority than the interrupted process is now ready to dispatch and dispatches the process. The time to switch context from a lower-priority process to a higher-priority process is included in the dispatch latency time.

Figure 11-3 illustrates the internal dispatch latency and application response time of a system. The response time is defined in terms of the amount of time a system takes to respond to an internal event. The dispatch latency of an internal event represents the amount of time that a process needs to wake up a higher priority process. The dispatch latency also includes the time that the system takes to dispatch the higher priority process.

The application response time is the amount of time that a driver takes to: wake up a higher-priority process, release resources from a low-priority process, reschedule the higher-priority task, calculate the response, and dispatch the task.

Interrupts can arrive and be processed during the dispatch latency interval. This processing increases the application response time, but is not attributed to the dispatch latency measurement. Therefore, this processing is not bounded by the dispatch latency guarantee.

FIGURE 11-3 Internal Dispatch Latency



With the new scheduling techniques provided with real-time SunOS, the system dispatch latency time is within specified bounds. As you can see in the following table, dispatch latency improves with a bounded number of processes.

TABLE 11-1 Real-time System Dispatch Latency

Workstation	Bounded Number of Processes	Arbitrary Number of Processes
SPARCstation 2	<0.5 milliseconds in a system with fewer than 16 active processes	1.0 milliseconds
SPARCstation 5	<0.3 millisecond	0.3 millisecond

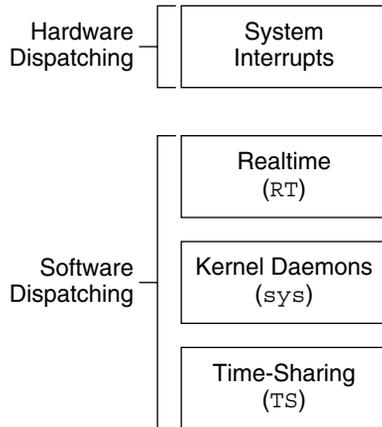
Scheduling Classes

The SunOS kernel dispatches processes by priority. The scheduler or dispatcher supports the concept of scheduling classes. Classes are defined as real-time (RT), system (SYS), and time-sharing (TS). Each class has a unique scheduling policy for dispatching processes within its class.

The kernel dispatches highest priority processes first. By default, real-time processes have precedence over sys and TS processes. Administrators can configure systems so that the priorities for TS processes and RT processes overlap.

The following figure illustrates the concept of classes as viewed by the SunOS kernel.

FIGURE 11-4 Dispatch Priorities for Scheduling Classes



Hardware interrupts, which cannot be controlled by software, have the highest priority. The routines that process interrupts are dispatched directly and immediately from interrupts, without regard to the priority of the current process.

Real-time processes have the highest default software priority. Processes in the RT class have a priority and *time quantum* value. RT processes are scheduled strictly on the basis of these parameters. As long as an RT process is ready to run, no SYS or TS process can run.

Fixed-priority scheduling enables critical processes to run in a predetermined order until completion. These priorities never change unless they are changed by an application.

An RT class process inherits the parent's time quantum, whether finite or infinite. A process with a finite time quantum runs until the time quantum expires. A process with a finite time quantum also stops running if the process blocks while waiting for an I/O event or is preempted by a higher-priority runnable real-time process. A process with an infinite time quantum ceases execution only when the process terminates, blocks, or is preempted.

The SYS class exists to schedule the execution of special system processes, such as paging, STREAMS, and the swapper. You cannot change the class of a process to the SYS class. The SYS class of processes has fixed priorities established by the kernel when the processes are started.

The time-sharing (TS) processes have the lowest priority. TS class processes are scheduled dynamically, with a few hundred milliseconds for each time slice. The TS scheduler switches context in round-robin fashion often enough to give every process an equal opportunity to run, depending upon:

- The time slice value
- The process history, which records when the process was last put to sleep

- Considerations for CPU utilization

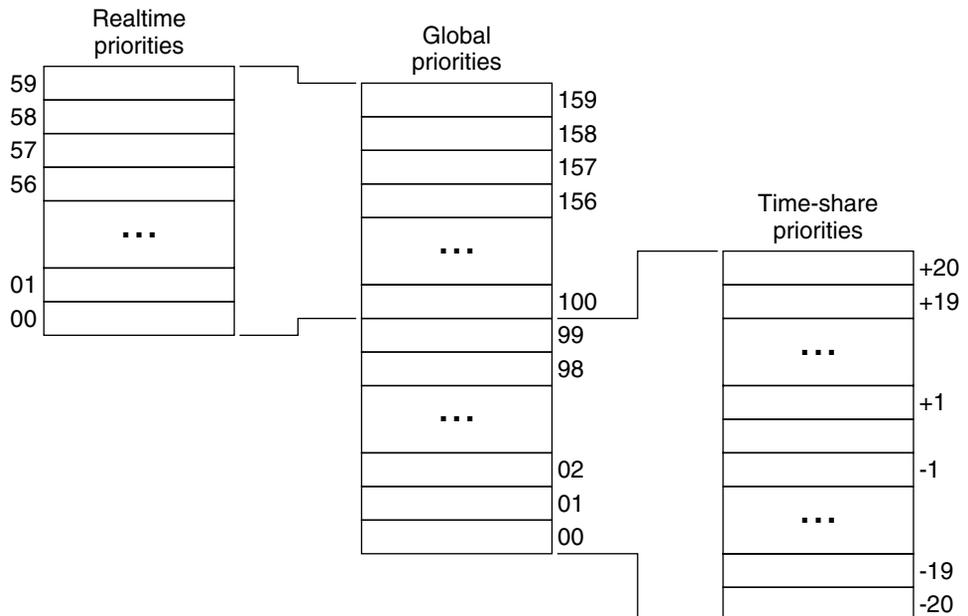
Default time-sharing policy gives larger time slices to processes with lower priority.

A child process inherits the scheduling class and attributes of the parent process through `fork(2)`. A process's scheduling class and attributes are unchanged by `exec(2)`.

Different algorithms dispatch each scheduling class. Class-dependent routines are called by the kernel to make decisions about CPU process scheduling. The kernel is class-independent, and takes the highest priority process off its queue. Each class is responsible for calculating a process's priority value for its class. This value is placed into the dispatch priority variable of that process.

As the following figure illustrates, each class algorithm has its own method of nominating the highest priority process to place on the global run queue.

FIGURE 11-5 Kernel Dispatch Queue



Each class has a set of priority levels that apply to processes in that class. A class-specific mapping maps these priorities into a set of global priorities. A set of global scheduling priority maps is not required to start with zero or be contiguous.

By default, the global priority values for time-sharing (TS) processes range from -20 to +20. These global priority values are mapped into the kernel from 0-40, with temporary assignments

as high as 99. The default priorities for real-time (RT) processes range from 0-59, and are mapped into the kernel from 100 to 159. The kernel's class-independent code runs the process with the highest global priority on the queue.

Dispatch Queue

The dispatch queue is a linear-linked list of processes with the same global priority. Each process has class-specific information attached to the process upon invocation. A process is dispatched from the kernel dispatch table in an order that is based on the process' global priority.

Dispatching Processes

When a process is dispatched, the context of the process is mapped into memory along with its memory management information, its registers, and its stack. Execution begins after the context mapping is done. Memory management information is in the form of hardware registers that contain the data that is needed to perform virtual memory translations for the currently running process.

Process Preemption

When a higher priority process becomes dispatchable, the kernel interrupts its computation and forces the context switch, preempting the currently running process. A process can be preempted at any time if the kernel finds that a higher-priority process is now dispatchable.

For example, suppose that process A performs a read from a peripheral device. Process A is put into the sleep state by the kernel. The kernel then finds that a lower-priority process B is runnable. Process B is dispatched and begins execution. Eventually, the peripheral device sends an interrupt, and the driver of the device is entered. The device driver makes process A runnable and returns. Rather than returning to the interrupted process B, the kernel now preempts B from processing, resuming execution of the awakened process A.

Another interesting situation occurs when several processes contend for kernel resources. A high-priority real-time process might be waiting for a resource held by a low-priority process. When the low-priority process releases the resource, the kernel preempts that process to resume execution of the higher-priority process.

Kernel Priority Inversion

Priority inversion occurs when a higher-priority process is blocked by one or more lower-priority processes for a long time. The use of synchronization primitives such as mutual-exclusion locks in the SunOS kernel can lead to priority inversion.

A process is *blocked* when the process must wait for one or more processes to relinquish resources. Prolonged blocking can lead to missed deadlines, even for low levels of utilization.

The problem of priority inversion has been addressed for mutual-exclusion locks for the SunOS kernel by implementing a basic priority inheritance policy. The policy states that a lower-priority process inherits the priority of a higher-priority process when the lower-priority process blocks the execution of the higher-priority process. This inheritance places an upper bound on the amount of time a process can remain blocked. The policy is a property of the kernel's behavior, not a solution that a programmer institutes through system calls or interface execution. User-level processes can still exhibit priority inversion, however.

User Priority Inversion

The issue of user priority inversion, and the means to deal with priority inversion, are discussed in “Mutual Exclusion Lock Attributes” in *Multithreaded Programming Guide*.

Interface Calls That Control Scheduling

The following interface calls control process scheduling.

Using `prioctl`

Control over scheduling of active classes is done with `prioctl(2)`. Class attributes are inherited through `fork(2)` and `exec(2)`, along with scheduling parameters and permissions required for priority control. This inheritance happens with both the RT and the TS classes.

`prioctl(2)` is the interface for specifying a real-time process, a set of processes, or a class to which the system call applies. `prioctlset(2)` also provides the more general interface for specifying an entire set of processes to which the system call applies.

The command arguments of `prioctl(2)` can be one of: `PC_GETCID`, `PC_GETCLINFO`, `PC_GETPARMS`, or `PC_SETPARMS`. The real or effective ID of the calling process must match the real or effective ID of the affected processes, or must have superuser privilege.

<code>PC_GETCID</code>	This command takes the name field of a structure that contains a recognizable class name. The class ID and an array of class attribute data are returned.
<code>PC_GETCLINFO</code>	This command takes the ID field of a structure that contains a recognizable class identifier. The class name and an array of class attribute data are returned.
<code>PC_GETPARMS</code>	This command returns the scheduling class identifier or the class specific scheduling parameters of one of the specified processes. Even though <code>idtype</code> and <code>id</code> might specify a big set, <code>PC_GETPARMS</code> returns the parameter of only one process. The class selects the process.
<code>PC_SETPARMS</code>	This command sets the scheduling class or the class-specific scheduling parameters of the specified process or processes.

Other interface calls

<code>sched_get_priority_max</code>	Returns the maximum values for the specified policy.
<code>sched_get_priority_min</code>	Returns the minimum values for the specified policy. For more information, see the <code>sched_get_priority_max(3R)</code> man page.
<code>sched_rr_get_interval</code>	Updates the specified timespec structure to the current execution time limit.
<code>sched_setparam</code> , <code>sched_getparam</code>	Sets or gets the scheduling parameters of the specified process.
<code>sched_yield</code>	Blocks the calling process until the calling process returns to the head of the process list.

Utilities That Control Scheduling

The administrative utilities that control process scheduling are `dispadmin(1M)` and `priocntl(1)`. Both of these utilities support the `priocntl(2)` system call with compatible options and loadable modules. These utilities provide system administration functions that control real-time process scheduling during runtime.

`priocntl(1)`

The `priocntl(1)` command sets and retrieves scheduler parameters for processes.

`dispadmin(1M)`

The `dispadmin(1M)` utility displays all current process scheduling classes by including the `-l` command line option during runtime. Process scheduling can also be changed for the class specified after the `-c` option, using `RT` as the argument for the real-time class.

The class options for `dispadmin(1M)` are in the following list:

- `-l` Lists scheduler classes currently configured
- `-c` Specifies the class with parameters to be displayed or to be changed
- `-g` Gets the dispatch parameters for the specified class
- `-r` Used with `-g`, specifies time quantum resolution
- `-s` Specifies a file where values can be located

A class-specific file that contains the dispatch parameters can also be loaded during runtime. Use this file to establish a new set of priorities that replace the default values that were

established during boot time. This class-specific file must assert the arguments in the format used by the `-g` option. Parameters for the RT class are found in the `rt_dptbl(4)`, and are listed in [Example 11-1](#).

To add an RT class file to the system, the following modules must be present:

- An `rt_init()` routine in the class module that loads the `rt_dptbl(4)`.
- An `rt_dptbl(4)` module that provides the dispatch parameters and a routine to return pointers to `config_rt_dptbl`.
- The `dispadmin(1M)` executable.

The following steps install a RT class dispatch table:

1. Load the class-specific module with the following command, where *module_name* is the class-specific module.

```
# modload /kernel/sched/module_name
```

2. Invoke the `dispadmin` command.

```
# dispadmin -c RT -s file_name
```

The file must describe a table with the same number of entries as the table that is being overwritten.

Configuring Scheduling

Associated with both scheduling classes is a parameter table, `rt_dptbl(4)`, and `ts_dptbl(4)`. These tables are configurable by using a loadable module at boot time, or with `dispadmin(1M)` during runtime.

Dispatcher Parameter Table

The in-core table for real-time establishes the properties for RT scheduling. The `rt_dptbl(4)` structure consists of an array of parameters, `struct rt_dpent_t`. Each of the *n* priority levels has one parameter. The properties of a given priority level are specified by the *i*th parameter structure in the array, `rt_dptbl[i]`.

A parameter structure consists of the following members, which are also described in the `/usr/include/sys/rt.h` header file.

- | | |
|-------------------------|--|
| <code>rt_globpri</code> | The global scheduling priority associated with this priority level. The <code>rt_globpri</code> values cannot be changed with <code>dispadmin(1M)</code> . |
| <code>rt_quantum</code> | The length of the time quantum allocated to processes at this level in ticks. For more information, see “Timestamp Interfaces” on page 273 . The time quantum value is only a default or starting value for processes at a particular level. The |

time quantum of a real-time process can be changed by using the `prctl(1)` command or the `prctl(2)` system call.

Reconfiguring `config_rt_dptbl`

A real-time administrator can change the behavior of the real-time portion of the scheduler by reconfiguring the `config_rt_dptbl` at any time. One method is described in the `rt_dptbl(4)` man page, in the section titled “Replacing the `rt_dptbl` Loadable Module.”

A second method for examining or modifying the real-time parameter table on a running system is through the `dispadm(1M)` command. Invoking `dispadm(1M)` for the real-time class enables retrieval of the current `rt_quantum` values in the current `config_rt_dptbl` configuration from the kernel's in-core table. When overwriting the current in-core table, the configuration file used for input to `dispadm(1M)` must conform to the specific format described in the `rt_dptbl(4)` man page.

Following is an example of prioritized processes `rt_dptent_t` with their associated time quantum `config_rt_dptbl[]` value as the processes might appear in `config_rt_dptbl[]`.

EXAMPLE 11-1 RT Class Dispatch Parameters

```
rt_dptent_t config_rt_dptbl[] = {
    /* prilevel Time quantum */
    100, 100, 129, 60, 130, 40,
    101, 100, 131, 40,
    102, 100, 132, 40,
    103, 100, 133, 40,
    104, 100, 134, 40,
    105, 100, 135, 40,
    106, 100, 136, 40,
    107, 100, 137, 40,
    108, 100, 138, 40,
    109, 100, 139, 40,
    110, 80, 140, 20,
    111, 80, 141, 20,
    112, 80, 142, 20,
    113, 80, 143, 20,
    114, 80, 144, 20,
    115, 80, 145, 20,
    116, 80, 146, 20,
    117, 80, 147, 20,
    118, 80, 148, 20,
    119, 80, 149, 20,
    120, 60, 150, 10,
    121, 60, 151, 10,
    122, 60, 152, 10,
    123, 60, 153, 10,
    124, 60, 154, 10,
    125, 60, 155, 10,
    126, 60, 156, 10,
    126, 60, 157, 10,
    127, 60, 158, 10,
    128, 60, 159, 10,
}
```

Memory Locking

Locking memory is one of the most important issues for real-time applications. In a real-time environment, a process must be able to guarantee continuous memory residence to reduce latency and to prevent paging and swapping.

This section describes the memory locking mechanisms that are available to real-time applications in SunOS.

Under SunOS, the memory residency of a process is determined by its current state, the total available physical memory, the number of active processes, and the processes' demand for memory. This residency is appropriate in a time-share environment. This residency is often unacceptable for a real-time process. In a real-time environment, a process must guarantee a memory residence to reduce the process' memory access and dispatch latency.

Real-time memory locking in SunOS is provided by a set of library routines. These routines allow a process running with superuser privileges to lock specified portions of its virtual address space into physical memory. Pages locked in this manner are exempt from paging until the pages are unlocked or the process exits.

The operating system has a system-wide limit on the number of pages that can be locked at any time. This limit is a tunable parameter whose default value is calculated at boot time. The default value is based on the number of page frames minus another percentage, currently set at ten percent.

Locking a Page

A call to `mlock(3C)` requests that one segment of memory be locked into the system's physical memory. The pages that make up the specified segment are faulted in. The lock count of each page is incremented. Any page whose lock count value is greater than zero is exempt from paging activity.

A particular page can be locked multiple times by multiple processes through different mappings. If two different processes lock the same page, the page remains locked until both processes remove their locks. However, within a given mapping, page locks do not nest. Multiple calls of locking interfaces on the same address by the same process are removed by a single unlock request.

If the mapping through which a lock has been performed is removed, the memory segment is implicitly unlocked. When a page is deleted through closing or truncating the file, the page is also implicitly unlocked.

Locks are not inherited by a child process after a `fork(2)` call. If a process that has some memory locked forks a child, the child must perform a memory locking operation on its own behalf to

lock its own pages. Otherwise, the child process incurs copy-on-write page faults, which are the usual penalties that are associated with forking a process.

Unlocking a Page

To unlock a page of memory, a process requests the release of a segment of locked virtual pages by a calling `munlock(3C)`. `munlock` decrements the lock counts of the specified physical pages. After decrementing a page's lock count to 0, the page swaps normally.

Locking All Pages

A superuser process can request that all mappings within its address space be locked by a call to `mlockall(3C)`. If the flag `MCL_CURRENT` is set, all the existing memory mappings are locked. If the flag `MCL_FUTURE` is set, every mapping that is added to an existing mapping or that replaces an existing mapping is locked into memory.

Recovering Sticky Locks

A page is permanently locked into memory when its lock count reaches 65535 (`0xFFFF`). The value `0xFFFF` is defined by implementation. This value might change in future releases. Pages that are locked in this manner cannot be unlocked. Reboot the system to recover.

High Performance I/O

This section describes I/O with real-time processes. In SunOS, the libraries supply two sets of interfaces and calls to perform fast, asynchronous I/O operations. The POSIX asynchronous I/O interfaces are the most recent standard. The SunOS environment also provides file and in-memory synchronization operations and modes to prevent information loss and data inconsistency.

Standard UNIX I/O is synchronous to the application programmer. An application that calls `read(2)` or `write(2)` usually waits until the system call has finished.

Real-time applications need asynchronous, bounded I/O behavior. A process that issues an asynchronous I/O call proceeds without waiting for the I/O operation to complete. The caller is notified when the I/O operation has finished.

Asynchronous I/O can be used with any SunOS file. Files are opened synchronously and no special flagging is required. An asynchronous I/O transfer has three elements: call, request, and

operation. The application calls an asynchronous I/O interface, the request for the I/O is placed on a queue, and the call returns immediately. At some point, the system dequeues the request and initiates the I/O operation.

Asynchronous and standard I/O requests can be intermingled on any file descriptor. The system maintains no particular sequence of read and write requests. The system arbitrarily resequences all pending read and write requests. If a specific sequence is required for the application, the application must insure the completion of prior operations before issuing the dependent requests.

POSIX Asynchronous I/O

POSIX asynchronous I/O is performed using `aiocb` structures. An `aiocb` control block identifies each asynchronous I/O request and contains all of the controlling information. A control block can be used for only one request at a time. A control block can be reused after its request has been completed.

A typical POSIX asynchronous I/O operation is initiated by a call to `aio_read` or `aio_write`. Either polling or signals can be used to determine the completion of an operation. If signals are used for completing operations, each operation can be uniquely tagged. The tag is then returned in the `si_value` component of the generated signal. See the [siginfo\(3HEAD\)](#) man page.

<code>aio_read</code>	Is called with an asynchronous I/O control block to initiate a read operation.
<code>aio_write</code>	Is called with an asynchronous I/O control block to initiate a write operation.
<code>aio_return</code> , <code>aio_error</code>	Are called to obtain return and error values, respectively, after an operation is known to have completed.
<code>aio_cancel</code>	Is called with an asynchronous I/O control block to cancel pending operations. <code>aio_cancel</code> can be used to cancel a specific request, if a request is specified by the control block. <code>aio_cancel</code> can also cancel all of the requests that are pending for the specified file descriptor.
<code>aio_fsync</code>	Queues an asynchronous <code>fsync</code> or <code>fdatasync</code> request for all of the pending I/O operations on the specified file.
<code>aio_suspend</code>	Suspends the caller as though one or more of the preceding asynchronous I/O requests had been made synchronously.

Oracle Solaris Asynchronous I/O

This section discusses asynchronous I/O operations in the Oracle Solaris operating environment.

Notification (SIGIO)

When an asynchronous I/O call returns successfully, the I/O operation has only been queued and waits to be done. The actual operation has a return value and a potential error identifier. This return value and potential error identifier would have been returned to the caller if the call had been synchronous. When the I/O is finished, both the return and error values are stored at a location given by the user at the time of the request as a pointer to an `aio_result_t`. The structure of the `aio_result_t` is defined in `<sys/asynch.h>`:

```
typedef struct aio_result_t {
    ssize_t    aio_return; /* return value of read or write */
    int       aio_errno;  /* errno generated by the IO */
} aio_result_t;
```

When the `aio_result_t` has been updated, a SIGIO signal is delivered to the process that made the I/O request.

Note that a process with two or more asynchronous I/O operations pending has no certain way to determine the cause of the SIGIO signal. A process that receives a SIGIO should check all its conditions that could be generating the SIGIO signal.

Using `aioread`

This command routine is the asynchronous version of `read(2)`. In addition to the normal read arguments, `aioread` takes the arguments that specify a file position and the address of an `aio_result_t` structure. The resulting information about the operation is stored in the `aio_result_t` structure. The file position specifies a seek to be performed within the file before the operation. Whether the `aioread` command call succeeds or fails, the file pointer is updated.

Using `aiowrite`

The `aiowrite` command routine is the asynchronous version of `write(2)`. In addition to the normal write arguments, `aiowrite` command takes arguments that specify a file position and the address of an `aio_result_t` structure. The resulting information about the operation is stored in the `aio_result_t` structure.

The file position specifies that a seek operation is to be performed within the file before the operation. If the command call succeeds, the file pointer is updated to the position that would have resulted in a successful seek and write. The file pointer is also updated when a write fails to allow for subsequent write requests.

Using `aiocancel`

This command routine attempts to cancel the asynchronous request whose `aio_result_t` structure is given as an argument. An `aiocancel` call succeeds only if the request is still queued. If the operation is in progress, `aiocancel` fails.

Using `aiowait`

A call to `aiowait` blocks the calling process until at least one outstanding asynchronous I/O operation is completed. The timeout parameter points to a maximum interval to wait for I/O completion. A timeout value of zero specifies that no wait is wanted. The `aiowait` command returns a pointer to the `aio_result_t` structure for the completed operation.

Using `poll()`

To determine the completion of an asynchronous I/O event synchronously rather than depend on a SIGIO interrupt, use `poll(2)`. You can also poll to determine the origin of a SIGIO interrupt.

`poll(2)` is slow when used on very large numbers of files. This problem is resolved by `poll(7d)`.

Using the `poll` Driver

Using `/dev/poll` provides a highly scalable way of polling a large number of file descriptors. This scalability is provided through a new set of APIs and a new driver, `/dev/poll`. The `/dev/poll` API is an alternative to, not a replacement of, `poll(2)`. Use `poll(7d)` to provide details and examples of the `/dev/poll` API. When used properly, the `/dev/poll` API scales much better than `poll(2)`. This API is especially suited for applications that satisfy the following criteria:

- Applications that repeatedly poll a large number of file descriptors
- Polled file descriptors that are relatively stable, meaning that the descriptors are not constantly closed and reopened
- The set of file descriptors that actually have polled events pending is small, comparing to the total number of file descriptors that are being polled

Using `close`

Files are closed by calling `close(2)`. The call to `close(2)` cancels any outstanding asynchronous I/O request that can be closed. `close(2)` waits for an operation that cannot be cancelled. For more information, see “Using `aiocancel`” on page 269. When `close(2)` returns, no asynchronous I/O is pending for the file descriptor. Only asynchronous I/O requests queued to the specified file descriptor are cancelled when a file is closed. Any I/O pending requests for other file descriptors are not cancelled.

Synchronized I/O

Applications might need to guarantee that information has been written to stable storage, or that file updates are performed in a particular order. Synchronized I/O provides for these needs.

Synchronization Modes

Under SunOS, a write operation succeeds when the system ensures that all written data is readable after any subsequent open of the file. This check assumes no failure of the physical storage medium. Data is successfully transferred for a read operation when an image of the data on the physical storage medium is available to the requesting process. An I/O operation is complete when the associated data has been successfully transferred, or when the operation has been diagnosed as unsuccessful.

An I/O operation has reached synchronized I/O data integrity completion when:

- For reads, the operation has been completed, or diagnosed if unsuccessful. The read is complete only when an image of the data has been successfully transferred to the requesting process. If the synchronized read operation is requested when pending write requests affect the data to be read, these write requests are successfully completed before the data is read.
- For writes, the operation has been completed, or diagnosed if unsuccessful. The write operation succeeds when the data specified in the write request is successfully transferred. Furthermore, all file system information required to retrieve the data must be successfully transferred.
- File attributes that are not necessary for data retrieval are not transferred prior to returning to the calling process.
- Synchronized I/O file integrity completion requires that all file attributes relative to the I/O operation be successfully transferred before returning to the calling process. Synchronized I/O file integrity completion is otherwise identical to synchronized I/O data integrity completion.

Synchronizing a File

`fsync(3C)` and `fdatasync` explicitly synchronize a file to secondary storage.

The `fsync(3C)` routine guarantees that the interface is synchronized at the I/O file integrity completion level. `fdatasync` guarantees that the interface is synchronized at level of I/O data integrity completion.

Applications can synchronize each I/O operation before the operation completes. Setting the `O_DSYNC` flag on the file description by using `open(2)` or `fcntl(2)` ensures that all I/O writes reach I/O data completion before the operation completes. Setting the `O_SYNC` flag on the file description ensures that all I/O writes have reached completion before the operation is indicated as completed. Setting the `O_RSYNC` flag on the file description ensures that all I/O reads

[read\(2\)](#) and `aioread` reach the same level of completion that is requested by the descriptor setting. The descriptor setting can be either `O_DSYNC` or `O_SYNC`.

Interprocess Communication

This section describes the interprocess communication (IPC) interfaces of SunOS as the interfaces relate to real-time processing. Signals, pipes, FIFOs, message queues, shared memory, file mapping, and semaphores are described here. For more information about the libraries, interfaces, and routines that are useful for interprocess communication, see [Chapter 6, “Interprocess Communication.”](#)

Processing Signals

The sender can use `sigqueue` to send a signal together with a small amount of information to a target process.

To queue subsequent occurrences of a pending signal, the target process must have the `SA_SIGINFO` bit set for the specified signal. See the [sigaction\(2\)](#) man page.

The target process normally receive signals asynchronously. To receive signals synchronously, block the signal and call either `sigwaitinfo` or `sigtimedwait`. See the [sigprocmask\(2\)](#) man page. This procedure causes the signal to be received synchronously. The value sent by the caller of `sigqueue` is stored in the `si_value` member of the `siginfo_t` argument. Leaving the signal unblocked causes the signal to be delivered to the signal handler specified by [sigaction\(2\)](#), with the value appearing in the `si_value` of the `siginfo_t` argument to the handler.

A specified number of signals with associated values can be sent by a process and remain undelivered. Storage for `{SIGQUEUE_MAX}` signals is allocated at the first call to `sigqueue`. Thereafter, a call to the command either successfully enqueues at the target process or fails within a bounded amount of time.

Pipes, Named Pipes, and Message Queues

Pipes, named pipes, and message queues behave similarly to character I/O devices. These interfaces have different methods of connecting. See [“Pipes Between Processes” on page 93](#) for more information about pipes. See [“Named Pipes” on page 95](#) for more information about named pipes. See [“System V Messages” on page 99](#) and [“POSIX Messages” on page 96](#) for more information about message queues.

Using Semaphores

Semaphores are also provided in both System V and POSIX styles. See [“System V Semaphores” on page 101](#) and [“POSIX Semaphores” on page 97](#) for more information.

Note that using semaphores can cause priority inversions unless priority inversions are explicitly avoided by the techniques mentioned earlier in this chapter.

Shared Memory

The fastest way for processes to communicate is directly, through a shared segment of memory. When more than two processes attempt to read and write shared memory simultaneously, the memory contents can become inaccurate. This potential inaccuracy is the major difficulty with using shared memory.

Asynchronous Network Communication

This section introduces asynchronous network communication, using sockets or Transport-Level Interface (TLI) for real-time applications. Asynchronous networking with sockets is done by setting an open socket, of type `SOCK_STREAM`, to asynchronous and non blocking. For more information on asynchronous sockets, see [“Advanced Socket Topics” on page 136](#). Asynchronous network processing of TLI events is supported using a combination of STREAMS asynchronous features and the non-blocking mode of the TLI library routines.

For more information on the Transport-Level Interface, see [Chapter 8, “Programming With XTI and TLI.”](#)

Modes of Networking

Both sockets and transport-level interface provide two modes of service: *connection-mode* and *connectionless-mode*.

Connection-mode service is circuit-oriented. This service enables the transmission of data over an established connection in a reliable, sequenced manner. This service also provides an identification procedure that avoids the overhead of address resolution and transmission during the data transfer phase. This service is attractive for applications that require relatively long-lived, datastream-oriented interactions.

Connectionless-mode service is message-oriented and supports data transfer in self-contained units with no logical relationship required among multiple units. A single service request passes all the information required to deliver a unit of data from the sender to the transport provider.

This service request includes the destination address and the data to be delivered. Connectionless-mode service is attractive for applications that involve short-term interactions that do not require guaranteed, in-sequence delivery of data. Connectionless transports are generally unreliable.

Timing Facilities

This section describes the timing facilities that are available for real-time applications under SunOS. Real-time applications that use these mechanisms require detailed information from the man pages of the routines that are listed in this section.

The timing interfaces of SunOS fall into two separate areas: *timestamps* and *interval timers*. The timestamp interfaces provide a measure of elapsed time. The timestamp interfaces also enable the application to measure the duration of a state or the time between events. Interval timers allow an application to wake up at specified times and to schedule activities based on the passage of time.

Timestamp Interfaces

Two interfaces provide timestamps. `gettimeofday(3C)` provides the current time in a *timeval* structure, representing the time in seconds and microseconds since midnight, Greenwich Mean Time, on January 1, 1970. `clock_gettime`, with a `clockid` of `CLOCK_REALTIME`, provides the current time in a *timespec* structure, representing in seconds and nanoseconds the same time interval returned by `gettimeofday(3C)`.

SunOS uses a hardware periodic timer. For some workstations, the hardware periodic timer is the sole source of timing information. If the hardware periodic timer is the sole source of timing information, the accuracy of timestamps is limited to the timer's resolution. For other platforms, a timer register with a resolution of one microsecond means that timestamps are accurate to one microsecond.

Interval Timer Interfaces

Real-time applications often schedule actions by using interval timers. Interval timers can be either of two types: a *one-shot* type or a *periodic* type.

A one-shot is an armed timer that is set to an expiration time relative to either a current time or an absolute time. The timer expires once and is disarmed. This type of a timer is useful for clearing buffers after the data has been transferred to storage, or to time-out an operation.

A periodic timer is armed with an initial expiration time, either absolute or relative, and a repetition interval. Every time the interval timer expires, the timer is reloaded with the repetition interval. The timer is then rearmed. This timer is useful for data logging or for servo-control. In calls to interval timer interfaces, time values that are smaller than the timer's resolution are rounded up to the next multiple of the hardware timer interval. This interval is typically 10ms.

SunOS has two sets of timer interfaces. The `setitimer(2)` and `getitimer(2)` interfaces operate fixed set timers, which are called the BSD timers, using the `timeval` structure to specify time intervals. The POSIX timers, which are created with the `timer_create` command, operate the POSIX clock, `CLOCK_REALTIME`. POSIX timer operations are expressed in terms of the `timespec` structure.

The `getitimer(2)` and `setitimer(2)` functions retrieve and establish, respectively, the value of the specified BSD interval timer. The three BSD interval timers that are available to a process include a real-time timer designated `ITIMER_REAL`. If a BSD timer is armed and allowed to expire, the system sends an appropriate signal to the process that set the timer.

The `timer_create` command routine can create up to `TIMER_MAX` POSIX timers. The caller can specify what signal and what associated value are sent to the process when the timer expires. The `timer_gettime` and `timer_gettime` routines retrieve and establish respectively the value of the specified POSIX interval timer. POSIX timers can expire while the required signal is pending delivery. The timer expirations are counted, and `timer_getoverrun` retrieves the count. The `timer_delete` command deallocates a POSIX timer.

The following example illustrates how to use `setitimer(2)` to generate a periodic interrupt, and how to control the arrival of timer interrupts.

EXAMPLE 11-2 Controlling Timer Interrupts

```
#include <unistd.h>
#include <signal.h>
#include <sys/time.h>

#define TIMERCNT 8

void timerhandler();
int timercnt;
struct timeval alarmtimes[TIMERCNT];

main()
{
    struct itimerval times;
    sigset_t sigset;
    int i, ret;
    struct sigaction act;
    siginfo_t si;

    /* block SIGALRM */
    sigemptyset (&sigset);
```

EXAMPLE 11-2 Controlling Timer Interrupts (Continued)

```

sigaddset (&sigset, SIGALRM);
sigprocmask (SIG_BLOCK, &sigset, NULL);

/* set up handler for SIGALRM */
act.sa_action = timerhandler;
sigemptyset (&act.sa_mask);
act.sa_flags = SA_SIGINFO;
sigaction (SIGALRM, &act, NULL);
/*
 * set up interval timer, starting in three seconds,
 * then every 1/3 second
 */
times.it_value.tv_sec = 3;
times.it_value.tv_usec = 0;
times.it_interval.tv_sec = 0;
times.it_interval.tv_usec = 333333;
ret = setitimer (ITIMER_REAL, &times, NULL);
printf ("main:setitimer ret = %d\n", ret);

/* now wait for the alarms */
sigemptyset (&sigset);
timerhandler (0, si, NULL);
while (timercnt < TIMERCNT) {
    ret = sigsuspend (&sigset);
}
printtimes();
}

void timerhandler (sig, siginfo, context)
int      sig;
siginfo_t *siginfo;
void     *context;
{
    printf ("timerhandler:start\n");
    gettimeofday (&alarmtimes[timercnt], NULL);
    timercnt++;
    printf ("timerhandler:timercnt = %d\n", timercnt);
}

printtimes ()
{
    int    i;

    for (i = 0; i < TIMERCNT; i++) {
        printf ("%ld.%016d\n", alarmtimes[i].tv_sec,
                alarmtimes[i].tv_usec);
    }
}

```


The Oracle Solaris ABI and ABI Tools

The Oracle Solaris Application Binary Interface (ABI) defines the interfaces that are available for the use of application developers. Conforming to the ABI enhances an application's binary stability. This chapter discusses the Oracle Solaris ABI and the tools provided to verify an application's compliance with the ABI, including:

- The definition and purpose of the Oracle Solaris ABI, discussed in “[Defining the Oracle Solaris ABI](#)” on page 278.
- The usage of the two ABI tools, `apptcert` and `appttrace`, discussed in “[Oracle Solaris ABI Tools](#)” on page 280.

What is the Oracle Solaris ABI?

The Oracle Solaris ABI is the set of supported run-time interfaces that are available for an application to use with the Oracle Solaris operating system. The most important components of the ABI are in the following list:

- The interfaces provided by the Oracle Solaris system libraries, which are documented in section 3 of the man pages
- The interfaces provided by the Oracle Solaris kernel system calls, which are documented in section 2 of the man pages
- The locations and formats of various system files and directories, which are documented in section 4 of the man pages
- The input and output syntax and semantics of Oracle Solaris utilities, which are documented in section 1 of the man pages

The main component of the Oracle Solaris ABI is the set of system library interfaces. The term *ABI* in this chapter refers only to that component. The ABI contains exclusively C language interfaces, as C is the only language for which the Oracle Solaris operating system provides interfaces.

C source code that is written to the Oracle Solaris API (Application Programming Interface) is transformed by the C compiler into a binary for one of four ABI versions. The versions are:

- 32-bit SPARC
- 64-bit SPARC
- 32-bit x86
- x64

While the ABI is very similar to the API, the source compilation process introduces several important differences:

- Compiler directives such as `#define` can alter or replace source-level constructs. The resulting binary might lack a symbol present in the source or include a symbol not present in the source.
- The compiler might generate processor-specific symbols, such as arithmetic instructions, which augment or replace source constructs.
- The compiler's binary layout might be specific to that compiler and the versions of the source language which the compiler accepts. In such cases, identical code compiled with different compilers might produce incompatible binaries.

For these reasons, source-level (API) compatibility does not provide a sufficient expectation of binary compatibility across Oracle Solaris releases.

The Oracle Solaris ABI is made up of the supported interfaces provided by the operating system. Some of the interfaces that are available in the system are intended for the exclusive use of the operating system. These exclusive interfaces are not available for use by an application. Prior to the SunOS 5.6 release, all of the interfaces in Oracle Solaris libraries were available for application developers to use. With the library symbol scoping technology available in the Oracle Solaris link editor, interfaces not intended for use outside of a library have their scope reduced to be purely local to the library. See the [Linker and Libraries Guide](#) for details. Due to system requirements, not all private interfaces can have such a reduced scope. These interfaces are labeled *private*, and are not included in the Oracle Solaris ABI.

Defining the Oracle Solaris ABI

The Oracle Solaris ABI is defined in the Oracle Solaris libraries. These definitions are provided by means of the library versioning technology and policies used in the link editor and run-time linker.

Symbol Versioning in Oracle Solaris Libraries

The Oracle Solaris link editor and run-time linker use two kinds of library versioning: file versioning and symbol versioning. In file versioning, a library is named with an appended

version number, such as `libc.so.1`. In a few cases, when an incompatible change is made to one or more public interfaces in that library, the version number is incremented (for example, to `libc.so.2`).

Library symbol versioning associates a set of symbols with a symbol version name. The following is an example mapfile for a hypothetical Solaris library, `libfoo.so.1`.

```
$mapfile_version 2
    SYMBOL_VERSION SUNWpublic {
        global:
            symbolA;
            symbolB;
            symbolC;
    };

    SYMBOL_VERSION SUNWprivate {
        global:
            __fooimpl;
        local:
            *;
    };
```

This mapfile indicates that `symbolA`, `symbolB`, and `symbolC` are associated with version `SUNWpublic`. The symbol `__fooimpl` is associated with `SUNWprivate`.

Note – The *local: ** directive in the mapfile causes any symbol in the library that is not explicitly associated with a named version, to be scoped locally to the library. Such locally scoped symbols are not visible outside the library. This convention ensures that symbols are only visible when associated with a symbol versioning name.

Using Symbol Versioning to Label the Oracle Solaris ABI

Since all visible symbols in a library belong to a named version, the naming scheme can be used to label the symbols' ABI status. This labeling is done by associating all private interfaces with a version name beginning with `SUNWprivate`. Public interfaces begin with other names, specifically:

- `SYSVABI`, for interfaces defined by the System V ABI definition
- `SISCD`, for interfaces defined by the SPARC International *SPARC Compliance Definition*
- `SUNWpublic`, for interfaces defined by Oracle Corporation
- `SUNW_x[y]`, for numbered public interfaces defined by Oracle in the older versions of the operating system.

The definition of the Oracle Solaris library ABI is therefore contained in the libraries, and consists of the set of symbols that are associated with symbol version names that do not begin with `SUNWprivate`. The `pvs` command lists the symbols in a library.

Oracle Solaris ABI Tools

The Oracle Solaris operating system provides two tools to verify that an application's use of Oracle Solaris interfaces conforms to the Oracle Solaris ABI. The `appcert` utility statically examines the Oracle Solaris library interfaces used by ELF binaries for instances of private interface usage. The `appcert` utility produces summary and detailed reports of any potential binary stability problems it finds. The `apptrace` tool uses the link-auditing capability of the run-time linker to dynamically trace Oracle Solaris library routine calls as the application runs. This capability enables developers to examine an application's use of the Oracle Solaris system interfaces.

The ABI tools enable easy, rapid identification of binaries that might have binary compatibility problems with a given Oracle Solaris release. To check binary stability, perform the following steps:

- **Use `appcert` on the current Oracle Solaris release for triage.** This identifies which binaries use problematic interfaces and which do not.
- **Use `apptrace` on the target Oracle Solaris release for verification.** This verifies whether interface compatibility problems exist by enabling dynamic observation of those interfaces as they are used.

appcert Utility

The `appcert` utility is a Perl script that statically examines ELF binaries and compares the library symbols used against a model of public interfaces and private interfaces in a given Oracle Solaris release. The utility runs on either SPARC or x86 platforms. The utility can check interface usage for both SPARC and x86 32-bit interfaces as well as the 64-bit interfaces on SPARC. Note that `appcert` only examines C language interfaces.

As new Oracle Solaris releases become available, some library interfaces might change their behavior or disappear entirely. These changes can affect the performance of applications that rely on those interfaces. The Oracle Solaris ABI defines runtime library interfaces that are safe and stable for application use. The `appcert` utility is designed to help developers verify an application's compliance with the Oracle Solaris ABI.

What appcert Checks

The `appcert` utility examines your applications for:

- Private symbol usage
- Static linking
- Unbound symbols

Private Symbol Usage

Private symbols are functions or data that is used by Oracle Solaris libraries to call each other. The semantic behavior of private symbols might change, and symbols might sometimes be removed. Such symbols are called *demoted symbols*. The mutable nature of private symbols introduces the potential for instability in applications that depend on private symbols.

Static Linking

The semantics of private symbol calls between Oracle Solaris libraries might change between releases. Therefore, the creation of static links to archives degrades an application's binary stability. Dynamic links to the archive's corresponding shared object file avoid this problem.

Unbound Symbols

The `appcert` utility uses the dynamic linker to resolve the library symbols that are used by the application being examined. Symbols that the dynamic linker cannot resolve are called *unbound symbols*. Unbound symbols might be caused by environment problems, such as an incorrectly set `LD_LIBRARY_PATH` variable. Unbound symbols might also be caused by build problems, such as omitting the definitions of the `-lib` or `-z` switches at compile time. While these examples are minor, unbound symbols that are reported by `appcert` might indicate a more serious problem, such as a dependency on a private symbol that no longer exists.

What appcert Does Not Check

If the object file `appcert` is examining depends on libraries, those dependencies must be recorded in the object. To do so, be sure to use the compiler's `-l` switch when compiling the code. If the object file depends on other shared libraries, those libraries must be accessible through `LD_LIBRARY_PATH` or `RPATH` at the time you run `appcert`.

The `appcert` application cannot check 64-bit applications unless the machine is running the 64-bit Solaris kernel. Since Solaris provides no 64-bit static libraries, `appcert` does not perform static-linking checks on 64-bit applications.

The `appcert` utility cannot examine:

- Object files that are completely or partially statically linked. A completely statically linked object is reported as unstable.
- Executable files that do not have the execute permission set. The `appcert` utility skips such executables. Shared objects without the execute permission set are examined normally.
- Object files whose user ID is set to `root`.
- Non-ELF executables, such as shell scripts.
- Oracle Solaris interfaces in languages other than C. The code need not be in C, but the call to the Oracle Solaris library must be.

Working with appcert

To check your application with `appcert`, type:

```
appcert object|directory
```

replacing *object|directory* with either:

- The complete list of objects you want `appcert` to examine
- The complete list of directories that contain such objects

Note – You might run `appcert` in an environment that is different from the environment in which the application runs. If these environments are different, `appcert` might not be able to correctly resolve references to Oracle Solaris library interfaces.

The `appcert` utility uses the Oracle Solaris runtime linker to construct a profile of interface dependencies for each executable or shared object file. This profile is used to determine the Oracle Solaris system interfaces upon which the application depends. The dependencies that are outlined in the profile are compared to the Oracle Solaris ABI to verify conformance. No private interfaces should be found.

The `appcert` utility recursively searches directories for object files, ignoring non-ELF object files. After `appcert` has finished checking the application, `appcert` prints a rollup report to the standard output, usually the screen. A copy of this report is written in the working directory, which is usually `/tmp/appcert.pid`, in a file that is named `Report`. In the subdirectory name, *pid* represents the 1–to–6 digit process ID of that particular instantiation of `appcert`. See “[appcert Results](#)” on [page 284](#) for more on the directory structure to which `appcert` writes output files.

appcert Options

The following options modify the behavior of the `appcert` utility. You can type any of these options at the command line, after the `appcert` command but before the *object|directory* operand.

-B Runs `appcert` in batch mode.

In batch mode, the report produced by `appcert` contains one line for each binary checked.

A line that begins with `PASS` indicates the binary that is named in that line did not trigger any `appcert` warnings.

A line that begins with `FAIL` indicates problems were found in that binary.

A line that begins with `INC` indicates the binary that is named in that line could not be completely checked.

- f *infile* The file *infile* should contain a list of files to check, with one file name per line. These files are added to any files already specified at the command line. If you use this switch, you do not need to specify an object or directory at the command line.
- h Prints usage information for `appcert`.
- L By default, `appcert` notes any shared objects in an application, and appends the directories in which the shared objects reside to `LD_LIBRARY_PATH`. The `-L` switch disables this behavior.
- n By default, `appcert` follows symbolic links when `appcert` searches directories for binaries to check. The `-n` switch disables this behavior.
- S Appends the Oracle Solaris library directories `/usr/openwin/lib` and `/usr/dt/lib` to `LD_LIBRARY_PATH`.
- w *working_dir* Specifies a directory in which to run the library components. Temporary files are also created in the directory specified by this switch. If this switch is not specified, `appcert` uses the `/tmp` directory.

Using `appcert` for Application Triage

The `appcert` utility can be used to quickly and easily discern which applications in a given set have potential stability problems. The following table lists some common binary stability problems.

TABLE 12-1 Common Binary Stability Problems

Problem	Course of Action
Use of a private symbol that is known to change	Eliminate use of symbol immediately.
Use of a private symbol that has not changed yet	Application can still be run for now, but eliminate use of symbol as soon as practical.
Static linking of a library with a shared object counterpart	Use shared object counterpart instead.
Static linking of a library with no shared object counterpart	Convert <code>.a</code> file to <code>.so</code> file by using the command <code>ld -z alltext ract</code> if possible. Otherwise, continue to use static library until shared object is available.
Use of a private symbol for which no public equivalent is available	Contact Oracle and request a public interface.
Use of a symbol that is deprecated, or use of a symbol that is planned for removal	Application can still be run for now, but eliminate use of symbol as soon as practical.

TABLE 12-1 Common Binary Stability Problems (Continued)

Problem	Course of Action
Use of a public symbol that has changed	Recompile.

Potential stability problems caused by the use of private interfaces might not occur on a given release. The behavior of private interfaces does not always change between releases. To verify that a private interface's behavior has changed in the target release, use the `appt race` tool. Usage of `appt race` is discussed in [“Using `appt race` for Application Verification” on page 286](#).

apptert Results

The results of the `apptert` utility's analysis of an application's object files are written to several files that are located in the `apptert` utility's working directory, typically `/tmp`. The main subdirectory under the working directory is `apptert.pid`, where `pid` is the process ID for that instantiation of `apptert`. The `apptert` utility's results are written to the following files:

Index	Contains the mapping between checked binaries and the subdirectory in which <code>apptert</code> output specific to that binary is located.							
Report	Contains a copy of the rollup report that is displayed on <code>stdout</code> when <code>apptert</code> is run.							
Skipped	Contains a list of binaries that <code>apptert</code> was asked to check but was forced to skip, along with the reason each binary was skipped. These reasons are in the following list: <ul style="list-style-type: none"> ▪ File is not a binary object ▪ File cannot be read by the user ▪ File name contains metacharacters ▪ File does not have the execute bit set 							
<code>objects/object_name</code>	A separate subdirectory is under the <code>objects</code> subdirectory for each object examined by <code>apptert</code> . Each of these subdirectories contains the following files: <table border="0" style="margin-left: 2em;"> <tbody> <tr> <td><code>check.demoted.symbols</code></td> <td>Contains a list of symbols that <code>apptert</code> suspects are demoted Oracle Solaris symbols.</td> </tr> <tr> <td><code>check.dynamic.private</code></td> <td>Contains a list of private Oracle Solaris symbols to which the object is directly bound.</td> </tr> <tr> <td><code>check.dynamic.public</code></td> <td>Contains a list of public Oracle Solaris symbols to which the object is directly bound.</td> </tr> </tbody> </table>		<code>check.demoted.symbols</code>	Contains a list of symbols that <code>apptert</code> suspects are demoted Oracle Solaris symbols.	<code>check.dynamic.private</code>	Contains a list of private Oracle Solaris symbols to which the object is directly bound.	<code>check.dynamic.public</code>	Contains a list of public Oracle Solaris symbols to which the object is directly bound.
<code>check.demoted.symbols</code>	Contains a list of symbols that <code>apptert</code> suspects are demoted Oracle Solaris symbols.							
<code>check.dynamic.private</code>	Contains a list of private Oracle Solaris symbols to which the object is directly bound.							
<code>check.dynamic.public</code>	Contains a list of public Oracle Solaris symbols to which the object is directly bound.							

<code>check.dynamic.unbound</code>	Contains a list of symbols not bound by the dynamic linker when running <code>ldd -r</code> . Lines returned by <code>ldd</code> containing “file not found” are also included.
<code>summary.dynamic</code>	Contains a printer-formatted summary of dynamic bindings in the objects <code>appcert</code> examined, including tables of public and private symbols used from each Oracle Solaris library.

Returns one of four exit values.

- 0 No potential sources of binary instability were found by `appcert`.
- 1 The `appcert` utility did not run successfully.
- 2 Some of the objects checked by `appcert` have potential binary stability problems.
- 3 The `appcert` utility did not find any binary objects to check.

Correcting Problems Reported by `appcert`

- **Private Symbol Use** – An application that depends on private symbols might not run on a Oracle Solaris release different from the one in which it was developed. This phenomenon occurs because private symbols that occur in a given Oracle Solaris release might behave differently or not be present in another release. If `appcert` reports private symbol usage in your application, rewrite the application to avoid the use of private symbols.
- **Demoted Symbols** – Demoted symbols are functions or data variables in a Oracle Solaris library that have been removed or have been scoped locally in a later Oracle Solaris release. An application that directly calls such a symbol fails to run on a release whose libraries do not export that symbol.
- **Unbound Symbols** – Unbound symbols are library symbols that are referenced by the application that the dynamic linker was unable to resolve when called by `appcert`. While unbound symbols are not always an indicator of poor binary stability, unbound symbols might indicate a more serious problem, such as dependencies on demoted symbols.
- **Obsolete Library** – An obsolete library might be removed from the Oracle Solaris operating system in a future release. The `appcert` utility flags any use of such a library. Applications that depend on such a library might not function in a future release that does not feature the library. To avoid this problem, do not use interfaces from obsolete libraries.
- **Use of `sys_errlist` or `sys_nerr`** – The use of the `sys_errlist` and `sys_nerr` symbols might degrade binary stability. A reference might be made past the end of the `sys_errlist` array. To avoid this risk, use `strerror` instead.

- **Use of strong and weak symbols** – The strong symbols that are associated with weak symbols are reserved as private because their behavior might change in future Oracle Solaris releases. Applications should only directly reference weak symbols. An example of a strong symbol is `_socket`, which is associated with the weak symbol `socket`.

Using `aptrace` for Application Verification

The `aptrace` utility is a C program which dynamically traces calls to Oracle Solaris library routines as an application runs. `aptrace` works on either SPARC or x86 platforms. `aptrace` can trace interface calls for both SPARC and x86 32-bit interfaces, as well as the 64-bit interfaces on SPARC. As with `apccert`, `aptrace` only examines C language interfaces.

Application Verification

After using `apccert` to determine an application is at risk of binary instability, `aptrace` helps assess the degree of risk in each case. To determine an application's binary compatibility with a given release, verify the successful use of each interface used by the application with `aptrace`.

The `aptrace` utility can verify that an application is using public interfaces correctly. For example, an application that is using the `open()` to open the administrative file `/etc/passwd` directly should instead use the appropriate programmatic interfaces. This ability to inspect the usage of the Oracle Solaris ABI enables easy and rapid identification of potential interface problems.

Running `aptrace`

The `aptrace` utility does not require any modification of the application being traced. To use `aptrace`, type `aptrace`, followed by any desired options along with the command line used to run the application of interest. The `aptrace` utility works by using the link-auditing capability of the runtime linker to intercept the application's calls to Oracle Solaris library interfaces. The `aptrace` utility then traces the calls by printing the names and values of the call's arguments and return value. The tracing output can be on a single line or arranged across multiple lines for readability. Public interfaces are printed in human-readable form. Private interfaces are printed in hexadecimal.

The `aptrace` utility enables selective tracing of calls, both at the level of individual interfaces and the level of libraries. For example, `aptrace` can trace calls to `printf()` coming from `libc.so.1`, or a range of calls within a specific library. The `aptrace` utility can also verbosely trace user-specified calls. The specifications that dictate `aptrace` behavior are governed by a syntax that is consistent with the usage of `truss(1)`. The `-f` option directs `aptrace` to follow forked child processes. The `-o` option specifies an output file for `aptrace` results.

The `appt race` utility traces only library-level calls and is loaded into the running application process, gaining a performance increase over `truss`. With the exception of `printf`, `appt race` cannot trace calls to functions that accept variable argument lists or examine the stack or other caller information, for example, `setcontext`, `getcontext`, `setjmp`, `longjmp`, and `vfork`.

Interpreting `appt race` Output

The following examples contain sample `appt race` output from tracing a simple one-binary application, `ls`.

EXAMPLE 12-1 Default Tracing Behavior

```
% appt race ls /etc/passwd
ls      -> libc.so.1:atexit(func = 0xff3cb8f0) = 0x0
ls      -> libc.so.1:atexit(func = 0x129a4) = 0x0
ls      -> libc.so.1:getuid() = 0x32c3
ls      -> libc.so.1:time(tloc = 0x23918) = 0x3b2fe4ef
ls      -> libc.so.1:isatty(fildev = 0x1) = 0x1
ls      -> libc.so.1:ioctl(0x1, 0x540d, 0xffbfff7ac)
ls      -> libc.so.1:ioctl(0x1, 0x5468, 0x23908)
ls      -> libc.so.1:setlocale(category = 0x6, locale = "") = "C"
ls      -> libc.so.1:calloc(nelem = 0x1, elsize = 0x40) = 0x23cd0
ls      -> libc.so.1:lstat64(path = "/etc/passwd", buf = 0xffbfff6b0) = 0x0
ls      -> libc.so.1:acl(pathp = "/etc/passwd", cmd = 0x3, nentries = 0x0,
    acldata = 0x0) = 0x4
ls      -> libc.so.1:qsort(base = 0x23cd0, nel = 0x1, width = 0x40,
    compar = 0x12038)
ls      -> libc.so.1:sprintf(buf = 0x233d0, format = 0x12af8, ...) = 0
ls      -> libc.so.1:strlen(s = "") = 0x0
ls      -> libc.so.1:strlen(s = "/etc/passwd") = 0xb
ls      -> libc.so.1:sprintf(buf = 0x233d0, format = 0x12af8, ...) = 0
ls      -> libc.so.1:strlen(s = "") = 0x0
ls      -> libc.so.1:printf(format = 0x12ab8, ...) = 11
ls      -> libc.so.1:printf(/etc/passwd
format = 0x12abc, ...) = 1
ls      -> libc.so.1:exit(status = 0)
```

The previous example shows the default tracing behavior, tracing every library call on the command `ls /etc/passwd`. The `appt race` utility prints a line of output for every system call, indicating:

- The name of the call
- The library the call is in
- The arguments and return values of the call

The output from `ls` is mixed in with the `appt race` output.

EXAMPLE 12-2 Selective Tracing

```
% appt race -t \*printf ls /etc/passwd
ls      -> libc.so.1:sprintf(buf = 0x233d0, format = 0x12af8, ...) = 0
ls      -> libc.so.1:sprintf(buf = 0x233d0, format = 0x12af8, ...) = 0
```

EXAMPLE 12-2 Selective Tracing *(Continued)*

```
ls      -> libc.so.1:printf(format = 0x12ab8, ...) = 11
ls      -> libc.so.1:printf(/etc/passwd
format = 0x12abc, ...) = 1
```

The previous example shows how `appt race` can selectively trace calls with regular-expression syntax. In the example, calls to interfaces ending in `printf`, which include `sprintf`, are traced in the same `ls` command as before. Consequently, `appt race` only traces the `printf` and `sprintf` calls.

EXAMPLE 12-3 Verbose Tracing

```
% appt race -v sprintf ls /etc/passwd
ls      -> libc.so.1:sprintf(buf = 0x233d0, format = 0x12af8, ...) = 0
  buf = (char *) 0x233d0 ""
  format = (char *) 0x12af8 "%s%s%s"
ls      -> libc.so.1:sprintf(buf = 0x233d0, format = 0x12af8, ...) = 0
  buf = (char *) 0x233d0 ""
  format = (char *) 0x12af8 "%s%s%s"
/etc/passwd
```

The previous example shows the verbose tracing mode, where the arguments to `sprintf` are printed on multiple output lines for readability. At the end, `appt race` displays the output of the `ls` command.

UNIX Domain Sockets

UNIX domain sockets are named with UNIX paths. For example, a socket might be named `/tmp/foo`. UNIX domain sockets communicate only between processes on a single host. Sockets in the UNIX domain are not considered part of the network protocols because they can be used to communicate only between processes on a single host.

Socket types define the communication properties visible to a user. The Internet domain sockets provide access to the TCP/IP transport protocols. The Internet domain is identified by the value `AF_INET`. Sockets exchange data only with sockets in the same domain.

Creating Sockets

The `socket(3SOCKET)` call creates a socket in the specified family and of the specified type.

```
s = socket(family, type, protocol);
```

If the protocol is unspecified (a value of `0`), the system selects a protocol that supports the requested socket type. The socket handle (a file descriptor) is returned.

The family is specified by one of the constants defined in `sys/socket.h`. Constants named `AF_*` specify the address format to use in interpreting names.

The following creates a datagram socket for intramachine use:

```
s = socket(AF_UNIX, SOCK_DGRAM, 0);
```

Set the *protocol* argument to `0`, the default protocol, in most situations.

Local Name Binding

A socket is created with no name. A remote process has no way to refer to a socket until an address is bound to the socket. Communicating processes are connected through addresses. In the UNIX family, a connection is composed of (usually) one or two path names. UNIX family sockets need not always be bound to a name. If they are, bound, duplicate ordered sets such as local pathname or foreign pathname can never exist. The path names cannot refer to existing files.

The `bind(3SOCKET)` call enables a process to specify the local address of the socket. This creates the local pathname ordered set, while `connect(3SOCKET)` and `accept(3SOCKET)` complete a socket's association by fixing the remote half of the address. Use `bind(3SOCKET)` as follows:

```
bind (s, name, namelen);
```

The socket handle is `s`. The bound name is a byte string that is interpreted by the supporting protocols. UNIX family names contain a path name and a family. The example shows binding the name `/tmp/foo` to a UNIX family socket.

```
#include <sys/un.h>
...
struct sockaddr_un addr;
...
strncpy(addr.sun_path, "/tmp/foo", sizeof(addr.sun_path));
addr.sun_family = AF_UNIX;
bind (s, (struct sockaddr *) &addr,
      strlen(addr.sun_path) + sizeof (addr.sun_family));
```

When determining the size of an `AF_UNIX` socket address, null bytes are not counted, which is why you can use `strlen(3C)`.

The file name referred to in `addr.sun_path` is created as a socket in the system file name space. The caller must have write permission in the directory where `addr.sun_path` is created. The file should be deleted by the caller when it is no longer needed. Delete `AF_UNIX` sockets with `unlink(1M)`.

Establishing a Connection

Connection establishment is usually asymmetric. One process acts as the client and the other as the server. The server binds a socket to a well-known address associated with the service and blocks on its socket for a connect request. An unrelated process can then connect to the server. The client requests services from the server by initiating a connection to the server's socket. On the client side, the `connect(3SOCKET)` call initiates a connection. In the UNIX family, this might appear as:

```
struct sockaddr_un server;
    server.sun_family = AF_UNIX;
    ...
    connect(s, (struct sockaddr *)&server, strlen(server.sun_path)
        + sizeof (server.sun_family));
```

See [“Connection Errors” on page 115](#) for information on connection errors. [“Data Transfer” on page 115](#) tells you how to transfer data. [“Closing Sockets” on page 116](#) tells you how to close a socket.

Index

A

- ABI, *See* application binary interface
- ABI differences from API, 278
- accept, 113, 290
- API differences from ABI, 278
- appcert
 - limitations, 281
 - syntax, 282–283
- application binary interface (ABI), 277–278
 - defined, 278–279
 - tools, 280–288
 - appcert, 280
 - apptrace, 280
- apptrace, 286–288
- asynchronous I/O
 - behavior, 254
 - endpoint service, 200
 - guaranteeing buffer state, 254
 - listen for network connection, 202
 - making connection request, 202
 - notification of data arrival, 200
 - opening a file, 202–204
 - using structure, 254
- Asynchronous Safe, 190
- asynchronous socket, 138–139, 139
- atomic updates to semaphores, 102
- attribute, finding in an SDP session structure, 35–36

B

- bind, 113, 290

- blocking mode
 - defined, 260
 - finite time quantum, 258
 - priority inversion, 260
 - time-sharing process, 253
- brk(2), 19
- broadcast, sending message, 145

C

- calloc, 18
- chmod(1), 85
- class
 - definition, 257
 - priority queue, 259
 - scheduling algorithm, 259
 - scheduling priorities, 257–260
- Client-server model, 131
- close, 116
- connect, 113, 114, 125, 290
- connection-mode
 - asynchronous network service, 201–202
 - asynchronously connecting, 201
 - definition, 272
 - using asynchronous connection, 201–202
- connectionless mode, asynchronous network service, 200–201
- connectionless-mode, definition, 272
- context switch, preempting a process, 260
- creation flags, IPC, 98

D

- daemon, inetd, 143–144
- datagram
 - socket, 110, 123–127, 134
- debugging dynamic memory, 18
- /dev/zero, mapping, 16
- dispatch, priorities, 258
- dispatch latency, 255
 - under realtime, 255–261
- dispatch table
 - configuring, 263–264
 - kernel, 260
- dynamic memory
 - allocation, 17–18
 - debugging, 18

E

- EWOLDBLOCK, 138
- examples, library mapfile, 279

F

- F_GETLK, 89
- F_SETOWN fcntl, 139
- fcntl(2), 87
- file and record locking, 84
- file descriptor
 - passing to another process, 202
 - transferring, 202–204
- file system
 - contiguous, 254
 - opening dynamically, 202
- file versioning, 278
- files, lock, 84
- free, 18

G

- gethostbyaddr, 129
- gethostbyname, 129
- getpeername, 144

- getservbyname, 130
- getservbyport, 130
- getservent, 130

H

- handle
 - socket, 113, 290
- host name mapping, 129
- hostent structure, 129

I

- I/O, *See* asynchronous I/O or synchronous I/O

- inet_ntoa, 129
- inetd, 131, 143
- init(1M), scheduler properties, 51
- interfaces
 - advanced I/O, 83
 - basic I/O, 82
 - IPC, 93
 - list file system control, 84
 - terminal I/O, 90

Internet

- host name mapping, 129
- port numbers, 142
- well known address, 130, 131

Interprocess Communication (IPC)

- using messages, 271
- using named pipes, 271
- using pipes, 271
- using semaphores, 272
- using shared memory, 272

- ioctl, SIOCATMARK, 136

IPC (interprocess communication), 93

- creation flags, 98
- interfaces, 98
- messages, 99–101
- permissions, 98
- semaphores, 101–105
- shared memory, 105–108

- IPC_RMID, 100

- IPC_SET, 100

IPC_STAT, 100
 IPPORT_RESERVED, 142

K

kernel
 class independent, 259
 context switch, 260
 dispatch table, 260
 preempting current process, 260
 queue, 254

L

libnsl, 190
 lockf(3C), 89
 locking
 advisory, 86
 F_GETLK, 89
 finding locks, 89
 mandatory, 86
 memory in realtime, 265
 opening a file for, 87
 record, 88
 removing, 88–89
 setting, 88–89
 supported file systems, 86–90
 testing locks, 89
 with fcntl(2), 87
 ls(1), 85

M

malloc, 18
 mapped files, 15, 16
 media, finding in an SDP session structure, 36
 media format, finding in an SDP session structure, 36–37
 memalign, 18
 memory
 locking, 265
 locking a page, 265–266

memory (*Continued*)
 locking all pages, 266
 number of locked pages, 265
 sticky locks, 266
 unlocking a page, 266
 memory allocation, dynamic, 17–18
 memory management, 19
 brk, 19
 interfaces, 15–17
 mlock, 16
 mlockall, 17
 mmap, 15, 16
 mprotect, 19
 msync, 17
 munmap, 16
 sbrk, 19
 sysconf, 19
 messages, 99–101
 mlock, 16
 mlockall, 17
 mmap, 15, 16
 mprotect, 19
 MSG_DONTROUTE, 116
 MSG_OOB, 116
 MSG_PEEK, 116, 136
 msgget(), 99
 msqid, 99
 msync, 17
 multiple connect (TLI), 194
 multithread safe, 189, 243
 munmap, 16

N

name-to-address translation
 inet, 245
 nis.so, 244
 straddr.so, 245
 switch.so, 244
 tcipip.so, 244
 named pipe, FIFO, 271
 netdir_free, 246
 netdir_getbyaddr, 246
 netdir_getbyname, 246

netdir_options, 247
netdir_perror, 247
netdir_sperror, 247
netent structure, 129
network

- asynchronous connection, 199, 272
- asynchronous service, 200–201
- asynchronous transfers, 200–201
- asynchronous use, 200
- connection-mode service, 272–273
- connectionless-mode service, 272
- programming models for real-time, 199–200
- services under realtime, 272–273
- using STREAMS asynchronously, 199, 272
- using Transport-Level Interface (TLI), 199

networked applications, 11
nice(1), 51
nice(2), 51
nis.so, 244
non-blocking mode

- configuring endpoint connections, 201
- defined, 199
- endpoint bound to service address, 201
- network service, 200
- polling for notification, 200
- service requests, 200
- Transport-Level Interface (TLI), 199
- using `t_connect()`, 201

nonblocking sockets, 137–138

O

optmgmt, 205, 207, 208
Oracle Solaris Studio, 18
out-of-band data, 136

P

performance, scheduler effect on, 51
permissions, IPC, 98
poll, 193
pollfd structure, 195, 196

polling

- for a connection request, 201
- notification of data, 200
- using `poll(2)`, 200

port numbers for Internet, 142
port to service mapping, 130
porting from TLI to XTI, 190
prcntl(1), 49
priority inversion

- defined, 253
- synchronization, 260–261

priority queue, linear linked list, 260
process

- defined for realtime, 251–254
- dispatching, 260
- highest priority, 252
- preemption, 260
- residence in memory, 265
- runaway, 254
- scheduling for realtime, 258
- setting priorities, 262

process priority

- global, 44
- setting and retrieving, 49

protoent structure, 130

R

real-time, scheduler class, 46
realloc, 18
recvfrom, 125
removing record locks, 88–89
response time

- blocking processes, 253
- bounds to I/O, 252
- degrading, 252–254
- inheriting priority, 253
- servicing interrupts, 252
- sharing libraries, 253
- sticky locks, 254

reversing operations for semaphores, 102
rpcbind, 245
Run Time Checking (RTC), 18
rwho, 134

S

- sbrk, 19
- sbrk(2), 19
- scheduler, 43, 53
 - classes, 258
 - configuring, 263–264
 - effect on performance, 51
 - priority, 257
 - real-time, 255
 - real-time policy, 46
 - scheduling classes, 257
 - system policy, 46
 - time-sharing policy, 45
 - using system calls, 261–262
 - using utilities, 262–263
- scheduler, class, 46
- sdp_add_attribute, 32
- sdp_add_bandwidth, 31
- sdp_add_connection, 30
- sdp_add_email, 30
- sdp_add_information, 29
- sdp_add_key, 32
- sdp_add_media, 32–33
- sdp_add_name, 29
- sdp_add_origin, 29
- sdp_add_phone, 30
- sdp_add_repeat, 31
- sdp_add_time, 31
- sdp_add_uri, 29–30
- sdp_add_zone, 31–32
- sdp_clone_session, 40–41
- sdp_delete_all_field, 37
- sdp_delete_all_media_field, 37–38
- sdp_delete_attribute, 38
- sdp_delete_media, 38
- sdp_find_attribute, 35–36
- sdp_find_media, 36
- sdp_find_media_rtpmap, 36–37
- sdp_free_session, 38
- sdp_new_session, 28
- sdp_parse, 39–40
- SDP session structure
 - finding an attribute in, 35–36
 - finding media format in, 36–37
- SDP session structure (*Continued*)
 - finding media in, 36
- sdp_session_to_str, 41
- select, 121, 136
- semaphores, 101–105
 - arbitrary simultaneous updates, 102
 - atomic updates, 102
 - reversing operations and SEM_UNDO, 102
 - undo structure, 102
- semget(), 102
- semop(), 102
- send, 125
- servent structure, 130
- service to port mapping, 130
- Session Description Protocol API
 - API framework, 25–28
 - attribute field, 32
 - bandwidth field, 31
 - cloning a session, 40–41
 - connection field, 30
 - converting a session to string, 41
 - creating a new session structure, 28–35
 - deleting attributes, 38
 - deleting fields, 37
 - deleting media, 38
 - deleting media fields, 37–38
 - email field, 30
 - finding an attribute, 35–36
 - finding media, 36
 - finding media format, 36–37
 - freeing a session, 38
 - information field, 29
 - key field, 32
 - library functions, 28–41
 - media field, 32–33
 - name field, 29
 - origin field, 29
 - parsing a structure, 39–40
 - repeat field, 31
 - sdp_new_session, 28
 - searching the SDP session structure, 35–37
 - shutting down a session structure, 37–38
 - telephone field, 30
 - time field, 31

Session Description Protocol API (*Continued*)

- URI field, 29–30
- utility functions, 38–41
- zone field, 31–32
- setting record locks, 88–89
- shared memory, 105–108
- shmget(), 106
- shutdown, 116
- SIGIO, 139
- SIOCATMARK ioctl, 136
- SIOCGIFCONF ioctl, 145
- SIOCGIFFLAGS ioctl, 146
- SOCK_DGRAM, 110, 143
- SOCK_RAW, 113
- SOCK_STREAM, 110, 140, 143
- socket
 - address binding, 140–142
 - AF_INET
 - bind, 113
 - create, 113
 - getservbyname, 130
 - getservbyport, 130
 - getservent, 130
 - inet_ntoa, 129
 - socket, 289
 - AF_UNIX
 - bind, 113, 290
 - create, 289
 - delete, 290
 - asynchronous, 138–139, 139
 - close, 116
 - connect stream, 116–121
 - datagram, 110, 123–127, 134
 - handle, 113, 290
 - initiate connection, 114, 290
 - multiplexed, 121
 - nonblocking, 137–138
 - out-of-band data, 116, 136
 - select, 121, 136
 - selecting protocols, 140
 - SIOCGIFCONF ioctl, 145
 - SIOCGIFFLAGS ioctl, 146
 - SOCK_DGRAM
 - connect, 125

socket, SOCK_DGRAM (*Continued*)

- recvfrom, 125, 136
- send, 125
- SOCK_STREAM, 140
 - F_GETOWN fcntl, 139
 - F_SETOWN fcntl, 139
 - out-of-band, 136
 - SIGIO signal, 139
 - SIGURG signal, 139
 - TCP port, 131
 - UDP port, 131
- Solaris library symbol versioning, *See* symbol versioning
- stream
 - data, 136
 - socket, 110, 116
- switch.so, 244
- symbol versioning, 278
- synchronous I/O
 - blocking, 266
 - critical timing, 252
- sysconf, 19

T

- t_accept, 212
- t_alloc, 209, 211
- t_bind, 210, 211
- t_close, 206, 211
- t_connect, 212
- T_DATAXFER, 209
- t_error, 211
- t_free, 211
- t_getinfo, 209, 211
- t_getstate, 211
- t_listen, 193, 209, 212
- t_look, 211
- t_open, 193, 209, 211
- t_optmgmt, 211
- t_rcv, 212
- t_rcvconnect, 212
- t_rcvdis, 210, 212
- t_rcvrel, 210, 212
- t_rcvuderr, 210, 212

t_rcvv, 213
 t_rcvvudata, 213
 t_snd, 212
 t_snddis, 192, 212
 t_sndrel, 210, 212
 t_sndreldata, 213
 t_sndudata, 212
 t_sndv, 213
 t_sndvudata, 213
 t_sync, 211
 t_sysconf, 213
 t_unbind, 211
 TCP, port, 131
 tcpip.so, 244
 time-sharing

- scheduler class, 45
- scheduler parameter table, 46

 timers

- applications, 273
- for interval timing, 273
- timestamping, 273
- using one-shot, 273
- using periodic type, 273

 tirdwr, 212
 tiuser.h, 190
 TLI

- asynchronous mode, 193–194
- broadcast, 210
- incoming events, 205–206
- multiple connection requests, 194
- opaque addresses, 210
- outgoing events, 204–205
- privileged ports, 210
- protocol independence, 209–210
- queue connect requests, 195
- queue multiple requests, 195
- read/write interface, 191–193
- socket comparison, 210
- state transitions, 206–209
- states, 204

 Transport-Level Interface (TLI), asynchronous endpoint, 200

U

UDP, port, 131
 undo structure for semaphores, 102
 unlink, 290
 updates, atomic for semaphores, 102
 usage, appt race, 286–288
 user priority, 45

V

valloc, 18
 versioning

- file, 278
- symbol, 278

 virtual memory, 19

X

XTI, 190
 xti.h, 190
 XTI Interface, 213
 XTI Utility Interfaces, 213
 XTI variables, getting, 213

Z

zero, 16

