



man pages section 3: Threads and Realtime Library Functions

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A.

Part No: 816-0216-06
December 2001

Copyright 2001 Sun Microsystems, Inc. 901 San Antonio Road Palo Alto, CA 94303-4900 U.S.A. All rights reserved.

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

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

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

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

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

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

Copyright 2001 Sun Microsystems, Inc. 901 San Antonio Road Palo Alto, CA 94303-4900 U.S.A. Tous droits réservés

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

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

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

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

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



011030@2471



Contents

Preface 9

Threads and Realtime Library Functions 15

aiocancel(3AIO)	16
aio_cancel(3RT)	17
aio_error(3RT)	19
aio_fsync(3RT)	21
aioread(3AIO)	23
aio_read(3RT)	25
aio_return(3RT)	28
aio_suspend(3RT)	29
aiowait(3AIO)	31
aio_write(3RT)	32
cancellation(3THR)	35
clock_settime(3RT)	41
cond_init(3THR)	43
condition(3THR)	48
door_bind(3DOOR)	50
door_call(3DOOR)	53
door_create(3DOOR)	56
door_cred(3DOOR)	58
door_info(3DOOR)	59
door_return(3DOOR)	61
door_revoke(3DOOR)	62
door_server_create(3DOOR)	63

fdatsync(3RT) 65
 libthread_db(3THR) 66
 lio_listio(3RT) 73
 mq_close(3RT) 77
 mq_getattr(3RT) 78
 mq_notify(3RT) 79
 mq_open(3RT) 81
 mq_receive(3RT) 84
 mq_send(3RT) 86
 mq_setattr(3RT) 88
 mq_unlink(3RT) 89
 mutex(3THR) 90
 mutex_init(3THR) 92
 nanosleep(3RT) 103
 proc_service(3PROC) 104
 ps_lgetregs(3PROC) 107
 ps_pglobal_lookup(3PROC) 109
 ps_pread(3PROC) 110
 ps_pstop(3PROC) 111
 pthread_atfork(3THR) 113
 pthread_attr_getdetachstate(3THR) 115
 pthread_attr_getguardsize(3THR) 116
 pthread_attr_getinheritsched(3THR) 118
 pthread_attr_getschedparam(3THR) 120
 pthread_attr_getschedpolicy(3THR) 121
 pthread_attr_getscope(3THR) 122
 pthread_attr_getstackaddr(3THR) 123
 pthread_attr_getstacksize(3THR) 124
 pthread_attr_init(3THR) 125
 pthread_cancel(3THR) 127
 pthread_cleanup_pop(3THR) 128
 pthread_cleanup_push(3THR) 129
 pthread_condattr_getpshared(3THR) 130
 pthread_condattr_init(3THR) 132
 pthread_cond_init(3THR) 134
 pthread_cond_signal(3THR) 136
 pthread_cond_wait(3THR) 138

pthread_create(3THR)	141
pthread_detach(3THR)	145
pthread_equal(3THR)	146
pthread_exit(3THR)	147
pthread_getconcurrency(3THR)	148
pthread_getschedparam(3THR)	150
pthread_getspecific(3THR)	152
pthread_join(3THR)	153
pthread_key_create(3THR)	155
pthread_key_delete(3THR)	157
pthread_kill(3THR)	158
pthread_mutexattr_getprioceiling(3THR)	159
pthread_mutexattr_getprotocol(3THR)	161
pthread_mutexattr_getpshared(3THR)	164
pthread_mutexattr_getrobust_np(3THR)	166
pthread_mutexattr_gettype(3THR)	168
pthread_mutexattr_init(3THR)	170
pthread_mutex_consistent_np(3THR)	171
pthread_mutex_getprioceiling(3THR)	173
pthread_mutex_init(3THR)	175
pthread_mutex_lock(3THR)	177
pthread_once(3THR)	180
pthread_rwlockattr_getpshared(3THR)	181
pthread_rwlockattr_init(3THR)	182
pthread_rwlock_init(3THR)	183
pthread_rwlock_rdlock(3THR)	185
pthread_rwlock_unlock(3THR)	187
pthread_rwlock_wrlock(3THR)	188
pthread_self(3THR)	190
pthread_setcancelstate(3THR)	191
pthread_setcanceltype(3THR)	193
pthread_sigmask(3THR)	195
pthread_testcancel(3THR)	200
rwlock(3THR)	201
schedctl_init(3SCHED)	204
sched_getparam(3RT)	206
sched_get_priority_max(3RT)	207

sched_getscheduler(3RT) 208
 sched_rr_get_interval(3RT) 209
 sched_setparam(3RT) 210
 sched_setscheduler(3RT) 212
 sched_yield(3RT) 214
 semaphore(3THR) 215
 sem_close(3RT) 219
 sem_destroy(3RT) 220
 sem_getvalue(3RT) 221
 sem_init(3RT) 222
 sem_open(3RT) 224
 sem_post(3RT) 227
 sem_unlink(3RT) 229
 sem_wait(3RT) 230
 shm_open(3RT) 233
 shm_unlink(3RT) 236
 sigqueue(3RT) 237
 sigwaitinfo(3RT) 239
 td_init(3THR) 241
 td_log(3THR) 242
 td_sync_get_info(3THR) 243
 td_ta_enable_stats(3THR) 247
 td_ta_event_addr(3THR) 249
 td_ta_get_nthreads(3THR) 253
 td_ta_map_addr2sync(3THR) 254
 td_ta_map_id2thr(3THR) 255
 td_ta_new(3THR) 256
 td_ta_setconcurrency(3THR) 258
 td_ta_sync_iter(3THR) 259
 td_thr_dbuspend(3THR) 261
 td_thr_getgregs(3THR) 262
 td_thr_get_info(3THR) 264
 td_thr_lockowner(3THR) 267
 td_thr_setprio(3THR) 268
 td_thr_setsigpending(3THR) 269
 td_thr_sleepinfo(3THR) 270
 td_thr_tsd(3THR) 271

td_thr_validate(3THR)	272
thr_create(3THR)	273
threads(3THR)	279
thr_exit(3THR)	286
thr_getconcurrency(3THR)	288
thr_getprio(3THR)	289
thr_join(3THR)	290
thr_keycreate(3THR)	292
thr_kill(3THR)	295
thr_main(3THR)	296
thr_min_stack(3THR)	297
thr_self(3THR)	298
thr_sigsetmask(3THR)	299
thr_stksegment(3THR)	304
thr_suspend(3THR)	305
thr_yield(3THR)	306
timer_create(3RT)	307
timer_delete(3RT)	309
timer_settime(3RT)	310

Index 313

Preface

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

Overview

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

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

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

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

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

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

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

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

Threads and Realtime Library Functions

aiocancel(3AIO)

NAME	aiocancel – cancel an asynchronous operation				
SYNOPSIS	<pre>cc [flag ...] file ... -laio [library ...] #include <sys/asynch.h> int aiocancel(aio_result_t *resultp);</pre>				
DESCRIPTION	<p>aiocancel() cancels the asynchronous operation associated with the result buffer pointed to by <i>resultp</i>. It may not be possible to immediately cancel an operation which is in progress and in this case, aiocancel() will not wait to cancel it.</p> <p>Upon successful completion, aiocancel() returns 0 and the requested operation is cancelled. The application will not receive the SIGIO completion signal for an asynchronous operation that is successfully cancelled.</p>				
RETURN VALUES	Upon successful completion, aiocancel() returns 0. Upon failure, aiocancel() returns -1 and sets errno to indicate the error.				
ERRORS	aiocancel() will fail if any of the following are true: EACCES The parameter <i>resultp</i> does not correspond to any outstanding asynchronous operation, although there is at least one currently outstanding. EFAULT <i>resultp</i> points to an address outside the address space of the requesting process. See NOTES. EINVAL There are not any outstanding requests to cancel.				
ATTRIBUTES	See attributes (5) for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
SEE ALSO	aioread(3AIO), aiowait(3AIO), attributes(5)				
NOTES	Passing an illegal address as <i>resultp</i> will result in setting errno to EFAULT <i>only</i> if it is detected by the application process.				

NAME	aio_cancel – cancel asynchronous I/O request
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <aio.h> int aio_cancel(int fildes, struct aiocb *aiocbp);</pre>
DESCRIPTION	<p>The <code>aio_cancel()</code> function attempts to cancel one or more asynchronous I/O requests currently outstanding against file descriptor <code>fildes</code>. The <code>aiocbp</code> argument points to the asynchronous I/O control block for a particular request to be canceled. If <code>aiocbp</code> is <code>NULL</code>, then all outstanding cancelable asynchronous I/O requests against <code>fildes</code> are canceled.</p> <p>Normal asynchronous notification occurs for asynchronous I/O operations that are successfully canceled. If there are requests that cannot be canceled, then the normal asynchronous completion process takes place for those requests when they are completed.</p> <p>For requested operations that are successfully canceled, the associated error status is set to <code>ECANCELED</code> and the return status is <code>-1</code>. For requested operations that are not successfully canceled, the <code>aiocbp</code> is not modified by <code>aio_cancel()</code>.</p> <p>If <code>aiocbp</code> is not <code>NULL</code>, then if <code>fildes</code> does not have the same value as the file descriptor with which the asynchronous operation was initiated, unspecified results occur.</p>
RETURN VALUES	<p>The <code>aio_cancel()</code> function returns the value <code>AIO_CANCELED</code> to the calling process if the requested operation(s) were canceled. The value <code>AIO_NOTCANCELED</code> is returned if at least one of the requested operation(s) cannot be canceled because it is in progress. In this case, the state of the other operations, if any, referenced in the call to <code>aio_cancel()</code> is not indicated by the return value of <code>aio_cancel()</code>. The application may determine the state of affairs for these operations by using <code>aio_error(3RT)</code>. The value <code>AIO_ALLDONE</code> is returned if all of the operations have already completed. Otherwise, the function returns <code>-1</code> and sets <code>errno</code> to indicate the error.</p>
ERRORS	<p>The <code>aio_cancel()</code> function will fail if:</p> <p><code>EBADF</code> The <code>fildes</code> argument is not a valid file descriptor.</p> <p><code>ENOSYS</code> The <code>aio_cancel()</code> function is not supported.</p>
USAGE	<p>The <code>aio_cancel()</code> function has a transitional interface for 64-bit file offsets. See <code>lf64(5)</code>.</p>
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p>

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

aio_cancel(3RT)

SEE ALSO aio_read(3RT), aio_return(3RT), attributes(5), aio(3HEAD), lf64(5),
signal(3HEAD)

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option.
Prior to this release, this function always returned -1 and set errno to ENOSYS.

NAME	aio_error – retrieve errors status for an asynchronous I/O operation
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <aio.h> int aio_error(const struct aiocb *aiocbp);</pre>
DESCRIPTION	The <code>aio_error()</code> function returns the error status associated with the <code>aiocb</code> structure referenced by the <code>aiocbp</code> argument. The error status for an asynchronous I/O operation is the <code>errno</code> value that would be set by the corresponding <code>read(2)</code> , <code>write(2)</code> , or <code>fsync(3C)</code> operation. If the operation has not yet completed, then the error status will be equal to <code>EINPROGRESS</code> .
RETURN VALUES	If the asynchronous I/O operation has completed successfully, then 0 is returned. If the asynchronous operation has completed unsuccessfully, then the error status, as described for <code>read(2)</code> , <code>write(2)</code> , and <code>fsync(3C)</code> , is returned. If the asynchronous I/O operation has not yet completed, then <code>EINPROGRESS</code> is returned.
ERRORS	<p>The <code>aio_error()</code> function will fail if:</p> <p><code>ENOSYS</code> The <code>aio_error()</code> function is not supported by the system.</p> <p>The <code>aio_error()</code> function may fail if:</p> <p><code>EINVAL</code> The <code>aiocbp</code> argument does not refer to an asynchronous operation whose return status has not yet been retrieved.</p>
USAGE	The <code>aio_error()</code> function has a transitional interface for 64-bit file offsets. See <code>lf64(5)</code> .
EXAMPLES	<p>EXAMPLE 1 The following is an example of an error handling routine using the <code>aio_error()</code> function.</p> <pre>#include <aio.h> #include <errno.h> #include <signal.h> struct aiocb my_aiocb; struct sigaction my_sigaction; void my_aio_handler(int, siginfo_t *, void *); . . . my_sigaction.sa_flags = SA_SIGINFO; my_sigaction.sa_sigaction = my_aio_handler; sigemptyset(&my_sigaction.sa_mask); (void) sigaction(SIGRTMIN, &my_sigaction, NULL); . . . my_aiocb.aio_sigevent.sigev_notify = SIGEV_SIGNAL; my_aiocb.aio_sigevent.sigev_signo = SIGRTMIN; my_aiocb.aio_sigevent.sigev_value.sival_ptr = &myaiocb; . . . (void) aio_read(&my_aiocb); . . . void my_aio_handler(int signo, siginfo_t *siginfo, void *context) { int my_errno; struct aiocb *my_aiocbp;</pre>

aio_error(3RT)

EXAMPLE 1 The following is an example of an error handling routine using the `aio_error()` function. (Continued)

```
my_aiocbp = siginfo->si_value.sival_ptr;
    if ((my_errno = aio_error(my_aiocb)) != EINPROGRESS) {
        int my_status = aio_return(my_aiocb);
        if (my_status >= 0) { /* start another operation */
            . . .
        } else { /* handle I/O error */
            . . .
        }
    }
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO `_exit(2)`, `close(2)`, `fork(2)`, `lseek(2)`, `read(2)`, `write(2)`, `aio(3HEAD)`, `aio_cancel(3RT)`, `aio_fsync(3RT)`, `aio_read(3RT)`, `aio_return(3RT)`, `aio_write(3RT)`, `lio_listio(3RT)`, `signal(3HEAD)`, `attributes(5)`, `lf64(5)`

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned `-1` and set `errno` to `ENOSYS`.

NAME	aio_fsync – asynchronous file synchronization
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <aio.h> int aio_fsync(int op, struct aiocb *aiocbp);</pre>
DESCRIPTION	<p>The <code>aio_fsync()</code> function asynchronously forces all I/O operations associated with the file indicated by the file descriptor <code>aio_fildes</code> member of the <code>aiocb</code> structure referenced by the <code>aiocbp</code> argument and queued at the time of the call to <code>aio_fsync()</code> to the synchronized I/O completion state. The function call returns when the synchronization request has been initiated or queued to the file or device (even when the data cannot be synchronized immediately).</p> <p>If <code>op</code> is <code>O_DSYNC</code>, all currently queued I/O operations are completed as if by a call to <code>fdatasync(3RT)</code>; that is, as defined for synchronized I/O data integrity completion. If <code>op</code> is <code>O_SYNC</code>, all currently queued I/O operations are completed as if by a call to <code>fsync(3C)</code>; that is, as defined for synchronized I/O file integrity completion. If the <code>aio_fsync()</code> function fails, or if the operation queued by <code>aio_fsync()</code> fails, then, as for <code>fsync(3C)</code> and <code>fdatasync(3RT)</code>, outstanding I/O operations are not guaranteed to have been completed.</p> <p>If <code>aio_fsync()</code> succeeds, then it is only the I/O that was queued at the time of the call to <code>aio_fsync()</code> that is guaranteed to be forced to the relevant completion state. The completion of subsequent I/O on the file descriptor is not guaranteed to be completed in a synchronized fashion.</p> <p>The <code>aiocbp</code> argument refers to an asynchronous I/O control block. The <code>aiocbp</code> value may be used as an argument to <code>aio_error(3RT)</code> and <code>aio_return(3RT)</code> in order to determine the error status and return status, respectively, of the asynchronous operation while it is proceeding. When the request is queued, the error status for the operation is <code>EINPROGRESS</code>. When all data has been successfully transferred, the error status will be reset to reflect the success or failure of the operation. If the operation does not complete successfully, the error status for the operation will be set to indicate the error. The <code>aio_sigevent</code> member determines the asynchronous notification to occur when all operations have achieved synchronized I/O completion. All other members of the structure referenced by <code>aiocbp</code> are ignored. If the control block referenced by <code>aiocbp</code> becomes an illegal address prior to asynchronous I/O completion, then the behavior is undefined.</p> <p>If the <code>aio_fsync()</code> function fails or the <code>aiocbp</code> indicates an error condition, data is not guaranteed to have been successfully transferred.</p> <p>If <code>aiocbp</code> is <code>NULL</code>, then no status is returned in <code>aiocbp</code>, and no signal is generated upon completion of the operation.</p>
RETURN VALUES	The <code>aio_fsync()</code> function returns 0 to the calling process if the I/O operation is successfully queued; otherwise, the function returns -1 and sets <code>errno</code> to indicate the error.

aio_fsync(3RT)

ERRORS The `aio_fsync()` function will fail if:

EAGAIN	The requested asynchronous operation was not queued due to temporary resource limitations.
EBADF	The <code>aio_fildes</code> member of the <code>aio_cb</code> structure referenced by the <code>aio_cbp</code> argument is not a valid file descriptor open for writing.
EINVAL	The system does not support synchronized I/O for this file.
EINVAL	A value of <code>op</code> other than <code>O_DSYNC</code> or <code>O_SYNC</code> was specified.
ENOSYS	The <code>aio_fsync()</code> function is not supported by the system.

In the event that any of the queued I/O operations fail, `aio_fsync()` returns the error condition defined for `read(2)` and `write(2)`. The error will be returned in the error status for the asynchronous `fsync(3C)` operation, which can be retrieved using `aio_error(3RT)`.

USAGE The `aio_fsync()` function has a transitional interface for 64-bit file offsets. See `lf64(5)`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `fcntl(2)`, `open(2)`, `read(2)`, `write(2)`, `aio_error(3RT)`, `aio_return(3RT)`, `fdatasync(3RT)`, `fsync(3C)`, `attributes(5)`, `fcntl(3HEAD)`, `aio(3HEAD)`, `lf64(5)`, `signal(3HEAD)`

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned `-1` and set `errno` to `ENOSYS`.

NAME	aio_read, aiowrite – read or write asynchronous I/O operations
SYNOPSIS	<pre>cc [flag ...] file ... -laio [library ...] #include <sys/types.h> #include <sys/asynch.h> int aio_read(int fildes, char *bufp, int bufs, off_t offset, int whence, aio_result_t *resultp); int aiowrite(int fildes, const char *bufp, int bufs, off_t offset, int whence, aio_result_t *resultp);</pre>
DESCRIPTION	<p><code>aio_read()</code> initiates one asynchronous <code>read(2)</code> and returns control to the calling program. The <code>read()</code> continues concurrently with other activity of the process. An attempt is made to read <code>bufs</code> bytes of data from the object referenced by the descriptor <code>fildes</code> into the buffer pointed to by <code>bufp</code>.</p> <p><code>aiowrite()</code> initiates one asynchronous <code>write(2)</code> and returns control to the calling program. The <code>write()</code> continues concurrently with other activity of the process. An attempt is made to write <code>bufs</code> bytes of data from the buffer pointed to by <code>bufp</code> to the object referenced by the descriptor <code>fildes</code>.</p> <p>On objects capable of seeking, the I/O operation starts at the position specified by <code>whence</code> and <code>offset</code>. These parameters have the same meaning as the corresponding parameters to the <code>lseek(2)</code> function. On objects not capable of seeking the I/O operation always start from the current position and the parameters <code>whence</code> and <code>offset</code> are ignored. The seek pointer for objects capable of seeking is not updated by <code>aio_read()</code> or <code>aiowrite()</code>. Sequential asynchronous operations on these devices must be managed by the application using the <code>whence</code> and <code>offset</code> parameters.</p> <p>The result of the asynchronous operation is stored in the structure pointed to by <code>resultp</code>:</p> <pre>int aio_return; /* return value of read() or write() */ int aio_errno; /* value of errno for read() or write() */</pre> <p>Upon completion of the operation both <code>aio_return</code> and <code>aio_errno</code> are set to reflect the result of the operation. <code>AIO_INPROGRESS</code> is not a value used by the system so the client may detect a change in state by initializing <code>aio_return</code> to this value.</p> <p>The application supplied buffer <code>bufp</code> should not be referenced by the application until after the operation has completed. While the operation is <i>in progress</i>, this buffer is in use by the operating system.</p> <p>Notification of the completion of an asynchronous I/O operation may be obtained synchronously through the <code>aio_wait(3AIO)</code> function, or asynchronously by installing a signal handler for the <code>SIGIO</code> signal. Asynchronous notification is accomplished by sending the process a <code>SIGIO</code> signal. If a signal handler is not installed for the <code>SIGIO</code> signal, asynchronous notification is disabled. The delivery of this instance of the <code>SIGIO</code> signal is reliable in that a signal delivered while the handler is executing is not</p>

aioread(3AIO)

lost. If the client ensures that `aiowait(3AIO)` returns nothing (using a polling timeout) before returning from the signal handler, no asynchronous I/O notifications are lost. The `aiowait(3AIO)` function is the only way to dequeue an asynchronous notification. Note: `SIGIO` may have several meanings simultaneously: for example, that a descriptor generated `SIGIO` and an asynchronous operation completed. Further, issuing an asynchronous request successfully guarantees that space exists to queue the completion notification.

`close(2)`, `exit(2)` and `execve()` (see `exec(2)`) will block until all pending asynchronous I/O operations can be canceled by the system.

It is an error to use the same result buffer in more than one outstanding request. These structures may only be reused after the system has completed the operation.

RETURN VALUES Upon successful completion, `aioread()` and `aiowrite()` return 0. Upon failure, `aioread()` and `aiowrite()` return -1 and set `errno` to indicate the error.

ERRORS `aioread()` and `aiowrite()` will fail if any of the following are true:

<code>EAGAIN</code>	The number of asynchronous requests that the system can handle at any one time has been exceeded
<code>EBADF</code>	<i>fdes</i> is not a valid file descriptor open for reading.
<code>EFAULT</code>	At least one of <i>bufp</i> points to an address outside the address space of the requesting process. See <code>NOTES</code> .
<code>EINVAL</code>	The parameter <i>resultp</i> is currently being used by an outstanding asynchronous request.
<code>EINVAL</code>	<i>offset</i> is not a valid offset for this file system type.
<code>ENOMEM</code>	Memory resources are unavailable to initiate request.

USAGE The `aioread()` and `aiowrite()` functions have transitional interfaces for 64-bit file offsets. See `lf64(5)`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `close(2)`, `exec(2)`, `exit(2)`, `llseek(2)`, `lseek(2)`, `open(2)`, `read(2)`, `write(2)`, `aiocancel(3AIO)`, `aiowait(3AIO)`, `sigvec(3UCB)`, `attributes(5)`, `lf64(5)`

NOTES Passing an illegal address to *bufp* will result in setting `errno` to `EFAULT` *only* if it is detected by the application process.

NAME	aio_read – asynchronous read from a file
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <aio.h> int aio_read(struct aiocb *aiocbp);</pre>
DESCRIPTION	<p>The <code>aio_read()</code> function allows the calling process to read <code>aiocbp->aio_nbytes</code> from the file associated with <code>aiocbp->aio_fildes</code> into the buffer pointed to by <code>aiocbp->aio_buf</code>. The function call returns when the read request has been initiated or queued to the file or device (even when the data cannot be delivered immediately). If <code>_POSIX_PRIORITIZED_IO</code> is defined and prioritized I/O is supported for this file, then the asynchronous operation is submitted at a priority equal to the scheduling priority of the process minus <code>aiocbp->aio_reqprio</code>. The <code>aiocbp</code> value may be used as an argument to <code>aio_error(3RT)</code> and <code>aio_return(3RT)</code> in order to determine the error status and return status, respectively, of the asynchronous operation while it is proceeding. If an error condition is encountered during queuing, the function call returns without having initiated or queued the request. The requested operation takes place at the absolute position in the file as given by <code>aio_offset</code>, as if <code>lseek(2)</code> were called immediately prior to the operation with an <code>offset</code> equal to <code>aio_offset</code> and a whence equal to <code>SEEK_SET</code>. After a successful call to enqueue an asynchronous I/O operation, the value of the file offset for the file is unspecified.</p> <p>The <code>aiocbp->aio_lio_opcode</code> field is ignored by <code>aio_read()</code>.</p> <p>The <code>aiocbp</code> argument points to an <code>aiocb</code> structure. If the buffer pointed to by <code>aiocbp->aio_buf</code> or the control block pointed to by <code>aiocbp</code> becomes an illegal address prior to asynchronous I/O completion, then the behavior is undefined.</p> <p>Simultaneous asynchronous operations using the same <code>aiocbp</code> produce undefined results.</p> <p>If <code>_POSIX_SYNCHRONIZED_IO</code> is defined and synchronized I/O is enabled on the file associated with <code>aiocbp->aio_fildes</code>, the behavior of this function is according to the definitions of synchronized I/O data integrity completion and synchronized I/O file integrity completion.</p> <p>For any system action that changes the process memory space while an asynchronous I/O is outstanding to the address range being changed, the result of that action is undefined.</p> <p>For regular files, no data transfer will occur past the offset maximum established in the open file description associated with <code>aiocbp->aio_fildes</code>.</p>
RETURN VALUES	The <code>aio_read()</code> function returns 0 to the calling process if the I/O operation is successfully queued; otherwise, the function returns -1 and sets <code>errno</code> to indicate the error.
ERRORS	The <code>aio_read()</code> function will fail if:

aio_read(3RT)

EAGAIN The requested asynchronous I/O operation was not queued due to system resource limitations.

ENOSYS The `aio_read()` function is not supported by the system.

Each of the following conditions may be detected synchronously at the time of the call to `aio_read()`, or asynchronously. If any of the conditions below are detected synchronously, the `aio_read()` function returns `-1` and sets `errno` to the corresponding value. If any of the conditions below are detected asynchronously, the return status of the asynchronous operation is set to `-1`, and the error status of the asynchronous operation will be set to the corresponding value.

EBADF The `aiocbp->aio_fildes` argument is not a valid file descriptor open for reading.

EINVAL The file offset value implied by `aiocbp->aio_offset` would be invalid, `aiocbp->aio_reqprio` is not a valid value, or `aiocbp->aio_nbytes` is an invalid value.

In the case that the `aio_read()` successfully queues the I/O operation but the operation is subsequently canceled or encounters an error, the return status of the asynchronous operation is one of the values normally returned by the `read(2)` function call. In addition, the error status of the asynchronous operation will be set to one of the error statuses normally set by the `read()` function call, or one of the following values:

EBADF The `aiocbp->aio_fildes` argument is not a valid file descriptor open for reading.

ECANCELED The requested I/O was canceled before the I/O completed due to an explicit `aio_cancel(3RT)` request.

EINVAL The file offset value implied by `aiocbp->aio_offset` would be invalid.

The following condition may be detected synchronously or asynchronously:

E_OVERFLOW The file is a regular file, `aiocbp->aio_nbytes` is greater than 0 and the starting offset in `aiocbp->aio_offset` is before the end-of-file and is at or beyond the offset maximum in the open file description associated with `aiocbp->aio_fildes`.

USAGE For portability, the application should set `aiocb->aio_reqprio` to 0.

The `aio_read()` function has a transitional interface for 64-bit file offsets. See `1f64(5)`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
----------------	-----------------

MT-Level	MT-Safe
----------	---------

SEE ALSO `close(2)`, `exec(2)`, `exit(2)`, `fork(2)`, `lseek(2)`, `read(2)`, `write(2)`, `aio_cancel(3RT)`, `aio_return(3RT)`, `lio_listio(3RT)`, `attributes(5)`, `aio(3HEAD)`, `lf64(5)`, `siginfo(3HEAD)`, `signal(3HEAD)`

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned `-1` and set `errno` to `ENOSYS`.

aio_return(3RT)

NAME	aio_return – retrieve return status of an asynchronous I/O operation				
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <aio.h> ssize_t aio_return(struct aiocb *aiocbp);</pre>				
DESCRIPTION	The <code>aio_return()</code> function returns the return status associated with the <code>aiocb</code> structure referenced by the <code>aiocbp</code> argument. The return status for an asynchronous I/O operation is the value that would be returned by the corresponding <code>read(2)</code> , <code>write(2)</code> , or <code>fsync(3C)</code> function call. If the error status for the operation is equal to <code>EINPROGRESS</code> , then the return status for the operation is undefined. The <code>aio_return()</code> function may be called exactly once to retrieve the return status of a given asynchronous operation; thereafter, if the same <code>aiocb</code> structure is used in a call to <code>aio_return()</code> or <code>aio_error(3RT)</code> , an error may be returned. When the <code>aiocb</code> structure referred to by <code>aiocbp</code> is used to submit another asynchronous operation, then <code>aio_return()</code> may be successfully used to retrieve the return status of that operation.				
RETURN VALUES	If the asynchronous I/O operation has completed, then the return status, as described for <code>read(2)</code> , <code>write(2)</code> , and <code>fsync(3C)</code> , is returned. If the asynchronous I/O operation has not yet completed, the results of <code>aio_return()</code> are undefined.				
ERRORS	The <code>aio_return()</code> function will fail if: EINVAL The <code>aiocbp</code> argument does not refer to an asynchronous operation whose return status has not yet been retrieved. ENOSYS The <code>aio_return()</code> function is not supported by the system.				
USAGE	The <code>aio_return()</code> function has a transitional interface for 64-bit file offsets. See <code>lf64(5)</code> .				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	<code>close(2)</code> , <code>exec(2)</code> , <code>exit(2)</code> , <code>fork(2)</code> , <code>lseek(2)</code> , <code>read(2)</code> , <code>write(2)</code> , <code>aio_cancel(3RT)</code> , <code>aio_fsync(3RT)</code> , <code>aio_read(3RT)</code> , <code>fsync(3C)</code> , <code>lio_listio(3RT)</code> , <code>attributes(5)</code> , <code>aio(3HEAD)</code> , <code>lf64(5)</code> , <code>signal(3HEAD)</code>				
NOTES	Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned <code>-1</code> and set <code>errno</code> to <code>ENOSYS</code> .				

NAME	aio_suspend – wait for asynchronous I/O request								
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <aio.h> int aio_suspend(const struct aiocb * const list[], int nent, const struct timespec *timeout);</pre>								
DESCRIPTION	<p>The <code>aio_suspend()</code> function suspends the calling thread until at least one of the asynchronous I/O operations referenced by the <code>list</code> argument has completed, until a signal interrupts the function, or, if <code>timeout</code> is not <code>NULL</code>, until the time interval specified by <code>timeout</code> has passed. If any of the <code>aiocb</code> structures in the list correspond to completed asynchronous I/O operations (that is, the error status for the operation is not equal to <code>EINPROGRESS</code>) at the time of the call, the function returns without suspending the calling thread. The <code>list</code> argument is an array of pointers to asynchronous I/O control blocks. The <code>nent</code> argument indicates the number of elements in the array. Each <code>aiocb</code> structure pointed to will have been used in initiating an asynchronous I/O request via <code>aio_read(3RT)</code>, <code>aio_write(3RT)</code>, or <code>lio_listio(3RT)</code>. This array may contain null pointers, which are ignored. If this array contains pointers that refer to <code>aiocb</code> structures that have not been used in submitting asynchronous I/O, the effect is undefined.</p> <p>If the time interval indicated in the <code>timespec</code> structure pointed to by <code>timeout</code> passes before any of the I/O operations referenced by <code>list</code> are completed, then <code>aio_suspend()</code> returns with an error.</p>								
RETURN VALUES	<p>If <code>aio_suspend()</code> returns after one or more asynchronous I/O operations have completed, it returns 0. Otherwise, it returns <code>-1</code>, and sets <code>errno</code> to indicate the error.</p> <p>The application may determine which asynchronous I/O completed by scanning the associated error and return status using <code>aio_error(3RT)</code> and <code>aio_return(3RT)</code>, respectively.</p>								
ERRORS	<p>The <code>aio_suspend()</code> function will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;"><code>EAGAIN</code></td> <td>No asynchronous I/O indicated in the list referenced by <code>list</code> completed in the time interval indicated by <code>timeout</code>.</td> </tr> <tr> <td style="vertical-align: top;"><code>EINTR</code></td> <td>A signal interrupted the <code>aio_suspend()</code> function. Note that, since each asynchronous I/O operation may possibly provoke a signal when it completes, this error return may be caused by the completion of one (or more) of the very I/O operations being awaited.</td> </tr> <tr> <td style="vertical-align: top;"><code>ENOMEM</code></td> <td>There is currently not enough available memory; the application can try again later.</td> </tr> <tr> <td style="vertical-align: top;"><code>ENOSYS</code></td> <td>The <code>aio_suspend()</code> function is not supported by the system.</td> </tr> </table>	<code>EAGAIN</code>	No asynchronous I/O indicated in the list referenced by <code>list</code> completed in the time interval indicated by <code>timeout</code> .	<code>EINTR</code>	A signal interrupted the <code>aio_suspend()</code> function. Note that, since each asynchronous I/O operation may possibly provoke a signal when it completes, this error return may be caused by the completion of one (or more) of the very I/O operations being awaited.	<code>ENOMEM</code>	There is currently not enough available memory; the application can try again later.	<code>ENOSYS</code>	The <code>aio_suspend()</code> function is not supported by the system.
<code>EAGAIN</code>	No asynchronous I/O indicated in the list referenced by <code>list</code> completed in the time interval indicated by <code>timeout</code> .								
<code>EINTR</code>	A signal interrupted the <code>aio_suspend()</code> function. Note that, since each asynchronous I/O operation may possibly provoke a signal when it completes, this error return may be caused by the completion of one (or more) of the very I/O operations being awaited.								
<code>ENOMEM</code>	There is currently not enough available memory; the application can try again later.								
<code>ENOSYS</code>	The <code>aio_suspend()</code> function is not supported by the system.								
USAGE	The <code>aio_suspend()</code> function has a transitional interface for 64-bit file offsets. See <code>lf64(5)</code> .								

aio_suspend(3RT)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO aio_fsync(3RT), aio_read(3RT), aio_return(3RT), aio_write(3RT),
lio_listio(3RT), attributes(5), aio(3HEAD), lf64(5), signal(3HEAD)

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option.
Prior to this release, this function always returned -1 and set errno to ENOSYS.

NAME	aiowait – wait for completion of asynchronous I/O operation				
SYNOPSIS	<pre>cc [flag ...] file ... -laio [library ...] #include <sys/async.h> #include <sys/time.h> aio_result_t *aiowait(const struct timeval *timeout);</pre>				
DESCRIPTION	<p>aiowait() suspends the calling process until one of its outstanding asynchronous I/O operations completes. This provides a synchronous method of notification.</p> <p>If <i>timeout</i> is a non-zero pointer, it specifies a maximum interval to wait for the completion of an asynchronous I/O operation. If <i>timeout</i> is a zero pointer, then aiowait() blocks indefinitely. To effect a poll, the <i>timeout</i> parameter should be non-zero, pointing to a zero-valued <i>timeval</i> structure.</p> <p>The <i>timeval</i> structure is defined in <sys/time.h> and contains the following members:</p> <pre>long tv_sec; /* seconds */ long tv_usec; /* and microseconds */</pre>				
RETURN VALUES	Upon successful completion, aiowait() returns a pointer to the result structure used when the completed asynchronous I/O operation was requested. Upon failure, aiowait() returns -1 and sets errno to indicate the error. aiowait() returns 0 if the time limit expires.				
ERRORS	<p>aiowait() will fail if any of the following are true:</p> <p>EFAULT <i>timeout</i> points to an address outside the address space of the requesting process. See NOTES.</p> <p>EINTR aiowait() was interrupted by a signal.</p> <p>EINVAL There are no outstanding asynchronous I/O requests.</p>				
ATTRIBUTES	See attributes (5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
SEE ALSO	aiocancel(3AIO), aioread(3AIO), attributes(5)				
NOTES	<p>aiowait() is the only way to dequeue an asynchronous notification. It may be used either inside a SIGIO signal handler or in the main program. One SIGIO signal may represent several queued events.</p> <p>Passing an illegal address as <i>timeout</i> will result in setting errno to EFAULT <i>only</i> if it is detected by the application process.</p>				

aio_write(3RT)

NAME	aio_write – asynchronous write to a file
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <aio.h> int aio_write(struct aiocb *aiocbp);</pre>
DESCRIPTION	<p>The <code>aio_write()</code> function allows the calling process to write <code>aiocbp->aio_nbytes</code> to the file associated with <code>aiocbp->aio_fildes</code> from the buffer pointed to by <code>aiocbp->aio_buf</code>. The function call returns when the write request has been initiated or, at a minimum, queued to the file or device. If <code>_POSIX_PRIORITIZED_IO</code> is defined and prioritized I/O is supported for this file, then the asynchronous operation is submitted at a priority equal to the scheduling priority of the process minus <code>aiocbp->aio_reqprio</code>. The <code>aiocbp</code> may be used as an argument to <code>aio_error(3RT)</code> and <code>aio_return(3RT)</code> in order to determine the error status and return status, respectively, of the asynchronous operation while it is proceeding.</p> <p>The <code>aiocbp</code> argument points to an <code>aiocb</code> structure. If the buffer pointed to by <code>aiocbp->aio_buf</code> or the control block pointed to by <code>aiocbp</code> becomes an illegal address prior to asynchronous I/O completion, then the behavior is undefined.</p> <p>If <code>O_APPEND</code> is not set for the file descriptor <code>aiocbp->aio_fildes</code>, then the requested operation takes place at the absolute position in the file as given by <code>aiocbp->aio_offset</code>, as if <code>lseek(2)</code> were called immediately prior to the operation with an <code>offset</code> equal to <code>aiocbp->aio_offset</code> and a <code>whence</code> equal to <code>SEEK_SET</code>. If <code>O_APPEND</code> is set for the file descriptor, write operations append to the file in the same order as the calls were made. After a successful call to enqueue an asynchronous I/O operation, the value of the file offset for the file is unspecified.</p> <p>The <code>aiocbp->aio_lio_opcode</code> field is ignored by <code>aio_write()</code>.</p> <p>Simultaneous asynchronous operations using the same <code>aiocbp</code> produce undefined results.</p> <p>If <code>_POSIX_SYNCHRONIZED_IO</code> is defined and synchronized I/O is enabled on the file associated with <code>aiocbp->aio_fildes</code>, the behavior of this function shall be according to the definitions of synchronized I/O data integrity completion and synchronized I/O file integrity completion.</p> <p>For any system action that changes the process memory space while an asynchronous I/O is outstanding to the address range being changed, the result of that action is undefined.</p> <p>For regular files, no data transfer will occur past the offset maximum established in the open file description associated with <code>aiocbp->aio_fildes</code>.</p>
RETURN VALUES	The <code>aio_write()</code> function returns 0 to the calling process if the I/O operation is successfully queued; otherwise, the function returns -1 and sets <code>errno</code> to indicate the error.

ERRORS The `aio_write()` function will fail if:

EAGAIN The requested asynchronous I/O operation was not queued due to system resource limitations.

ENOSYS The `aio_write()` function is not supported by the system.

Each of the following conditions may be detected synchronously at the time of the call to `aio_write()`, or asynchronously. If any of the conditions below are detected synchronously, the `aio_write()` function returns `-1` and sets `errno` to the corresponding value. If any of the conditions below are detected asynchronously, the return status of the asynchronous operation is set to `-1`, and the error status of the asynchronous operation will be set to the corresponding value.

EBADF The `aioctx->aio_fildes` argument is not a valid file descriptor open for writing.

EINVAL The file offset value implied by `aioctx->aio_offset` would be invalid, `aioctx->aio_reqprio` is not a valid value, or `aioctx->aio_nbytes` is an invalid value.

In the case that the `aio_write()` successfully queues the I/O operation, the return status of the asynchronous operation will be one of the values normally returned by the `write(2)` function call. If the operation is successfully queued but is subsequently canceled or encounters an error, the error status for the asynchronous operation contains one of the values normally set by the `write()` function call, or one of the following:

EBADF The `aioctx->aio_fildes` argument is not a valid file descriptor open for writing.

EINVAL The file offset value implied by `aioctx->aio_offset` would be invalid.

ECANCELED The requested I/O was canceled before the I/O completed due to an explicit `aio_cancel(3RT)` request.

The following condition may be detected synchronously or asynchronously:

EFBIG The file is a regular file, `aioctx->aio_nbytes` is greater than 0 and the starting offset in `aioctx->aio_offset` is at or beyond the offset maximum in the open file description associated with `aioctx->aio_fildes`.

USAGE The `aio_write()` function has a transitional interface for 64-bit file offsets. See `lf64(5)`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
----------------	-----------------

aio_write(3RT)

MT-Level	MT-Safe
----------	---------

SEE ALSO aio_cancel(3RT), aio_error(3RT), aio_read(3RT), aio_return(3RT),
lio_listio(3RT), close(2), _exit(2), fork(2), lseek(2), write(2),
attributes(5), aio(3HEAD), lf64(5), signal(3HEAD)

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option.
Prior to this release, this function always returned -1 and set errno to ENOSYS.

NAME	cancellation – overview of concepts related to POSIX thread cancellation															
DESCRIPTION	<table border="1"> <thead> <tr> <th data-bbox="459 394 846 428">FUNCTION</th> <th data-bbox="846 394 1433 428">ACTION</th> </tr> </thead> <tbody> <tr> <td data-bbox="459 428 846 462">pthread_cancel</td> <td data-bbox="846 428 1433 462">Cancels thread execution.</td> </tr> <tr> <td data-bbox="459 462 846 495">pthread_setcancelstate</td> <td data-bbox="846 462 1433 495">Sets the cancellation <i>state</i> of a thread.</td> </tr> <tr> <td data-bbox="459 495 846 529">pthread_setcanceltype</td> <td data-bbox="846 495 1433 529">Sets the cancellation <i>type</i> of a thread.</td> </tr> <tr> <td data-bbox="459 529 846 562">pthread_testcancel</td> <td data-bbox="846 529 1433 562">Creates a cancellation point in the calling thread.</td> </tr> <tr> <td data-bbox="459 562 846 596">pthread_cleanup_push</td> <td data-bbox="846 562 1433 596">Pushes a cleanup handler routine.</td> </tr> <tr> <td data-bbox="459 596 846 630">pthread_cleanup_pop</td> <td data-bbox="846 596 1433 630">Pops a cleanup handler routine.</td> </tr> </tbody> </table>		FUNCTION	ACTION	pthread_cancel	Cancels thread execution.	pthread_setcancelstate	Sets the cancellation <i>state</i> of a thread.	pthread_setcanceltype	Sets the cancellation <i>type</i> of a thread.	pthread_testcancel	Creates a cancellation point in the calling thread.	pthread_cleanup_push	Pushes a cleanup handler routine.	pthread_cleanup_pop	Pops a cleanup handler routine.
FUNCTION	ACTION															
pthread_cancel	Cancels thread execution.															
pthread_setcancelstate	Sets the cancellation <i>state</i> of a thread.															
pthread_setcanceltype	Sets the cancellation <i>type</i> of a thread.															
pthread_testcancel	Creates a cancellation point in the calling thread.															
pthread_cleanup_push	Pushes a cleanup handler routine.															
pthread_cleanup_pop	Pops a cleanup handler routine.															
Cancellation	<p>Thread cancellation allows a thread to terminate the execution of any application thread in the process. Cancellation is useful when further operations of one or more threads are undesirable or unnecessary.</p> <p>An example of a situation that could benefit from using cancellation is an asynchronously-generated cancel condition such as a user requesting to close or exit some running operation. Another example is the completion of a task undertaken by a number of threads, such as solving a maze. While many threads search for the solution, one of the threads might solve the puzzle while the others continue to operate. Since they are serving no purpose at that point, they should all be canceled.</p>															
Planning Steps	<p>Planning and programming for most cancellations follow this pattern:</p> <ol style="list-style-type: none"> 1. Identify which threads you want to cancel, and insert <code>pthread_cancel(3THR)</code> statements. 2. Identify system-defined cancellation points where a thread that might be canceled could have changed system or program state that should be restored. See the <code>Cancellation Points</code> for a list. 3. When a thread changes the system or program state just before a cancellation point, and should restore that state before the thread is canceled, place a cleanup handler before the cancellation point with <code>pthread_cleanup_push(3THR)</code>. Wherever a thread restores the changed state, pop the cleanup handler from the cleanup stack with <code>pthread_cleanup_pop(3THR)</code>. 4. Know whether the threads you are canceling call into cancel-unsafe libraries, and disable cancellation with <code>pthread_setcancelstate(3THR)</code> before the call into the library. See <code>Cancellation State</code> and <code>Cancel-Safe</code>. 5. To cancel a thread in a procedure that contains no cancellation points, insert your own cancellation points with <code>pthread_testcancel(3THR)</code>. <code>pthread_testcancel(3THR)</code> creates cancellation points by testing for pending cancellations and performing those cancellations if they are found. Push and pop cleanup handlers around the cancellation point, if necessary (see Step 3, above). 															

cancellation(3THR)

Cancellation Points

The system defines certain points at which cancellation can occur (cancellation points), and you can create additional cancellation points in your application with `pthread_testcancel(3THR)`.

The following cancellation points are defined by the system (system-defined cancellation points): `aio_suspend(3RT)`, `close(2)`, `creat(2)`, `getmsg(2)`, `getpmsg(2)`, `lockf(3C)`, `mq_receive(3RT)`, `mq_send(3RT)`, `msgrcv(2)`, `msgsnd(2)`, `msync(3C)`, `nanosleep(3RT)`, `open(2)`, `pause(2)`, `poll(2)`, `pread(2)`, `pthread_cond_timedwait(3THR)`, `pthread_cond_wait(3THR)`, `pthread_join(3THR)`, `pthread_testcancel(3THR)`, `putmsg(2)`, `putpmsg(2)`, `pwrite(2)`, `read(2)`, `readv(2)`, `select(3C)`, `sem_wait(3RT)`, `sigpause(3C)`, `sigwaitinfo(3RT)`, `sigsuspend(2)`, `sigtimedwait(3RT)`, `sigwait(2)`, `sleep(3C)`, `sync(2)`, `system(3C)`, `tcdrain(3C)`, `usleep(3C)`, `wait(2)`, `waitid(2)`, `waitpid(2)`, `wait3(3C)`, `write(2)`, `writew(2)`, and `fcntl(2)`, when specifying `F_SETLKW` as the command.

When cancellation is asynchronous, cancellation can occur at any time (before, during, or after the execution of the function defined as the cancellation point). When cancellation is deferred (the default case), cancellation occurs only within the scope of a function defined as a cancellation point (after the function is called and before the function returns). See `Cancellation Type` for more information about deferred and asynchronous cancellation.

Choosing where to place cancellation points and understanding how cancellation affects your program depend upon your understanding of both your application and of cancellation mechanics.

Typically, any call that might require a long wait should be a cancellation point. Operations need to check for pending cancellation requests when the operation is about to block indefinitely. This includes threads waiting in `pthread_cond_wait(3THR)` and `pthread_cond_timedwait(3THR)`, threads waiting for the termination of another thread in `pthread_join(3THR)`, and threads blocked on `sigwait(2)`.

A mutex is explicitly *not* a cancellation point and should be held for only the minimal essential time.

Most of the dangers in performing cancellations deal with properly restoring invariants and freeing shared resources. For example, a carelessly canceled thread might leave a mutex in a locked state, leading to a deadlock. Or it might leave a region of memory allocated with no way to identify it and therefore no way to free it.

Cleanup Handlers

When a thread is canceled, it should release resources and clean up the state that is shared with other threads. So, whenever a thread that might be canceled changes the state of the system or of the program, be sure to push a cleanup handler with `pthread_cleanup_push(3THR)` before the cancellation point.

When a thread is canceled, all the currently-stacked cleanup handlers are executed in last-in-first-out (LIFO) order. Each handler is run in the scope in which it was pushed.

When the last cleanup handler returns, the thread-specific data destructor functions are called. Thread execution terminates when the last destructor function returns.

When, in the normal course of the program, an uncanceled thread restores state that it had previously changed, be sure to pop the cleanup handler (that you had set up where the change took place) using `pthread_cleanup_pop(3THR)`. That way, if the thread is canceled later, only currently-changed state will be restored by the handlers that are left in the stack.

Be sure to pop the handler in the same scope in which it was pushed. Also, make sure that each push statement has a matching pop statement, or compiler errors will be generated.

Cancellation State

Most programmers will use only the default cancellation state of `PTHREAD_CANCEL_ENABLE`, but can choose to change the state by using `pthread_setcancelstate(3THR)`, which determines whether a thread is cancelable at all. With the default *state* of `PTHREAD_CANCEL_ENABLE`, cancellation is enabled, and the thread is cancelable at points determined by its cancellation *type*. See Cancellation Type.

If the *state* is `PTHREAD_CANCEL_DISABLE`, cancellation is disabled, the thread is not cancelable at any point, and all cancellation requests to it are held pending.

You might want to disable cancellation before a call to a cancel-unsafe library, restoring the old cancel state when the call returns from the library. See `Cancel-Safe` for explanations of cancel safety.

Cancellation Type

A thread's cancellation type is set with `pthread_setcanceltype(3THR)`, and determines whether the thread can be canceled anywhere in its execution, or only at cancellation points.

With the default *type* of `PTHREAD_CANCEL_DEFERRED`, the thread is cancelable only at cancellation points, and then only when cancellation is enabled.

If the *type* is `PTHREAD_CANCEL_ASYNCHRONOUS`, the thread is cancelable at any point in its execution (assuming, of course, that cancellation is enabled). Try to limit regions of asynchronous cancellation to sequences with no external dependencies that could result in dangling resources or unresolved state conditions. Using asynchronous cancellation is discouraged because of the danger involved in trying to guarantee correct cleanup handling at absolutely every point in the program.

Cancellation Type/State Table		
Type	State	
		Enabled (Default)

cancellation(3THR)

Cancellation Type/State Table		
Deferred (Default)	Cancellation occurs when the target thread reaches a cancellation point and a cancel is pending. (Default)	All cancellation requests to the target thread are held pending.
Asynchronous	Receipt of a <code>pthread_cancel(3T)</code> call causes immediate cancellation.	All cancellation requests to the target thread are held pending; as soon as cancellation is re-enabled, pending cancellations are executed immediately.

Cancel-Safe With the arrival of POSIX cancellation, the *cancel-safe* level has been added to the list of MT-Safety levels See `Intro(3)`. An application or library is cancel-safe whenever it has arranged for cleanup handlers to restore system or program state wherever cancellation can occur. The application or library is specifically *Deferred-cancel-safe* when it is cancel-safe for threads whose cancellation type is `PTHREAD_CANCEL_DEFERRED` See `Cancellation State`. It is specifically *Asynchronous-cancel-safe* when it is cancel-safe for threads whose cancellation type is `PTHREAD_CANCEL_ASYNC`.

Obviously, it is easier to arrange for deferred cancel safety, as this requires system and program state protection only around cancellation points. In general, expect that most applications and libraries are *not* Asynchronous-cancel-safe.

POSIX Threads Only Note: The cancellation functions described in this reference page are available for POSIX threads, only (the Solaris threads interfaces do not provide cancellation functions).

EXAMPLES **EXAMPLE 1** Cancellation example

The following short C++ example shows the pushing/popping of cancellation handlers, the disabling/enabling of cancellation, the use of `pthread_testcancel()`, and so on. The `free_res()` cancellation handler in this example is a dummy function that simply prints a message, but that would free resources in a real application. The function `f2()` is called from the main thread, and goes deep into its call stack by calling itself recursively.

Before `f2()` starts running, the newly created thread has probably posted a cancellation on the main thread since the main thread calls `thr_yield()` right after creating `thread2`. Because cancellation was initially disabled in the main thread, through a call to `pthread_setcancelstate()`, the call to `f2()` from `main()` continues and constructs `X` at each recursive call, even though the main thread has a pending cancellation.

When `f2()` is called for the fifty-first time (when "`i == 50`"), `f2()` enables cancellation by calling `pthread_setcancelstate()`. It then establishes a cancellation point for

EXAMPLE 1 Cancellation example (Continued)

itself by calling `pthread_testcancel()`. (Because a cancellation is pending, a call to a cancellation point such as `read(2)` or `write(2)` would also cancel the caller here.)

After the `main()` thread is canceled at the fifty-first iteration, all the cleanup handlers that were pushed are called in sequence; this is indicated by the calls to `free_res()` and the calls to the destructor for `X`. At each level, the C++ runtime calls the destructor for `X` and then the cancellation handler, `free_res()`. The print messages from `free_res()` and `X`'s destructor show the sequence of calls.

At the end, the main thread is joined by `thread2`. Because the main thread was canceled, its return status from `pthread_join()` is `PTHREAD_CANCELED`. After the status is printed, `thread2` returns, killing the process (since it is the last thread in the process).

```
#include <pthread.h>
#include <sched.h>
extern "C" void thr_yield(void);

extern "C" void printf(...);

struct X {
    int x;
    X(int i){x = i; printf("X(%d) constructed.\n", i);}
    ~X(){ printf("X(%d) destroyed.\n", x);}
};

void
free_res(void *i)
{
    printf("Freeing `%d'\n", i);
}

char* f2(int i)
{
    try {
        X dummy(i);
        pthread_cleanup_push(free_res, (void *)i);
        if (i == 50) {
            pthread_setcancelstate(PTHREAD_CANCEL_ENABLE, NULL);
            pthread_testcancel();
        }
        f2(i+1);
        pthread_cleanup_pop(0);
    }
    catch (int) {
        printf("Error: In handler.\n");
    }
    return "f2";
}

void *
thread2(void *tid)
```

cancellation(3THR)

EXAMPLE 1 Cancellation example *(Continued)*

```
{
    void *sts;

    printf("I am new thread :%d\n", pthread_self());

    pthread_cancel((pthread_t)tid);

    pthread_join((pthread_t)tid, &sts);

    printf("main thread cancelled due to %d\n", sts);

    return (sts);
}

main()
{
    pthread_setcancelstate(PTHREAD_CANCEL_DISABLE, NULL);
    pthread_create(NULL, NULL, thread2, (void *)pthread_self());
    thr_yield();
    printf("Returned from %s\n", f2(0));
}
```

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [read\(2\)](#), [sigwait\(2\)](#), [write\(2\)](#), [intro\(3\)](#), [condition\(3THR\)](#), [pthread_cleanup_pop\(3THR\)](#), [pthread_cleanup_push\(3THR\)](#), [pthread_exit\(3THR\)](#), [pthread_join\(3THR\)](#), [pthread_setcancelstate\(3THR\)](#), [pthread_setcanceltype\(3THR\)](#), [pthread_testcancel\(3THR\)](#), [setjmp\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	clock_settime, clock_gettime, clock_getres – high-resolution clock operations
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <time.h> int clock_settime(clockid_t clock_id, const struct timespec *tp); int clock_gettime(clockid_t clock_id, struct timespec *tp); int clock_getres(clockid_t clock_id, struct timespec *res);</pre>
DESCRIPTION	<p>The <code>clock_settime()</code> function sets the specified clock, <code>clock_id</code>, to the value specified by <code>tp</code>. Time values that are between two consecutive non-negative integer multiples of the resolution of the specified clock are truncated down to the smaller multiple of the resolution.</p> <p>The <code>clock_gettime()</code> function returns the current value <code>tp</code> for the specified clock, <code>clock_id</code>.</p> <p>The resolution of any clock can be obtained by calling <code>clock_getres()</code>. Clock resolutions are system-dependent and cannot be set by a process. If the argument <code>res</code> is not NULL, the resolution of the specified clock is stored in the location pointed to by <code>res</code>. If <code>res</code> is NULL, the clock resolution is not returned. If the time argument of <code>clock_settime()</code> is not a multiple of <code>res</code>, then the value is truncated to a multiple of <code>res</code>.</p> <p>A clock may be systemwide (that is, visible to all processes) or per-process (measuring time that is meaningful only within a process).</p> <p>A <code>clock_id</code> of <code>CLOCK_REALTIME</code> is defined in <code><time.h></code>. This clock represents the realtime clock for the system. For this clock, the values returned by <code>clock_gettime()</code> and specified by <code>clock_settime()</code> represent the amount of time (in seconds and nanoseconds) since the Epoch. Additional clocks may also be supported. The interpretation of time values for these clocks is unspecified.</p> <p>A <code>clock_id</code> of <code>CLOCK_HIGHRES</code> represents the non-adjustable, high-resolution clock for the system. For this clock, the value returned by <code>clock_gettime(3RT)</code> represents the amount of time (in seconds and nanoseconds) since some arbitrary time in the past; it is not correlated in any way to the time of day, and thus is not subject to resetting or drifting by way of <code>adjtime(2)</code>, <code>ntp_adjtime(2)</code>, <code>settimeofday(3C)</code>, or <code>clock_settime()</code>. The time source for this clock is the same as that for <code>gethrtime(3C)</code>.</p> <p>Additional clocks may also be supported. The interpretation of time values for these clocks is unspecified.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	The <code>clock_settime()</code> , <code>clock_gettime()</code> and <code>clock_getres()</code> functions will fail if:

clock_settime(3RT)

EINVAL The *clock_id* argument does not specify a known clock.
ENOSYS The functions `clock_settime()`, `clock_gettime()`, and `clock_getres()` are not supported by this implementation.

The `clock_settime()` function will fail if:

EINVAL The *tp* argument to `clock_settime()` is outside the range for the given clock ID; or the *tp* argument specified a nanosecond value less than zero or greater than or equal to 1000 million.

The `clock_settime()` function may fail if:

EPERM The requesting process does not have the appropriate privilege to set the specified clock.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>clock_gettime()</code> is Async-Signal-Safe

SEE ALSO `time(2)`, `ctime(3C)`, `gethrtime(3C)`, `time(3HEAD)`, `timer_gettime(3RT)`, `attributes(5)`

NAME	cond_init, cond_wait, cond_timedwait, cond_reltimedwait, cond_signal, cond_broadcast, cond_destroy – condition variables
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> #include <synch.h> int cond_init(cond_t *cvp, int type, void *arg); int cond_wait(cond_t *cvp, mutex_t *mp); int cond_timedwait(cond_t *cvp, mutex_t *mp, timestruc_t *abstime); int cond_reltimedwait(cond_t *cvp, mutex_t *mp, timestruc_t *reltime); int cond_signal(cond_t *cvp); int cond_broadcast(cond_t *cvp); int cond_destroy(cond_t *cvp);</pre>
Initialize	<p>Condition variables and mutexes should be global. Condition variables that are allocated in writable memory can synchronize threads among processes if they are shared by the cooperating processes (see <code>mmap(2)</code>) and are initialized for this purpose.</p> <p>The scope of a condition variable is either intra-process or inter-process. This is dependent upon whether the argument is passed implicitly or explicitly to the initialization of that condition variable. A condition variable does not need to be explicitly initialized. A condition variable is initialized with all zeros, by default, and its scope is set to within the calling process. For inter-process synchronization, a condition variable must be initialized once, and only once, before use.</p> <p>A condition variable must not be simultaneously initialized by multiple threads or re-initialized while in use by other threads.</p> <p>Attributes of condition variables can be set to the default or customized at initialization.</p> <p>The <code>cond_init()</code> function initializes the condition variable pointed to by <code>cvp</code>. A condition variable can have several different types of behavior, specified by <code>type</code>. No current type uses <code>arg</code> although a future type may specify additional behavior parameters with <code>arg</code>. The <code>type</code> argument can take one of the following values:</p> <p><code>USYNC_THREAD</code> The condition variable can synchronize threads only in this process. This is the default.</p> <p><code>USYNC_PROCESS</code> The condition variable can synchronize threads in this process and other processes. Only one process should initialize the condition variable. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not</p>

cond_init(3THR)

allocate it in such shared memory.

Initializing condition variables can also be accomplished by allocating in zeroed memory, in which case, a *type* of `USYNC_THREAD` is assumed.

If default condition variable attributes are used, statically allocated condition variables can be initialized by the macro `DEFAULTTCV`.

Default condition variable initialization (intra-process):

```
cond_t cvp;  
  
cond_init(&cvp, NULL, NULL); /*initialize condition variable with default*/
```

or

```
cond_init(&cvp, USYNC_THREAD, NULL);
```

or

```
cond_t cond = DEFAULTTCV;
```

Customized condition variable initialization (inter-process):

```
cond_init(&cvp, USYNC_PROCESS, NULL); /* initialize cv with  
inter-process scope */
```

Condition Wait

The condition wait interface allows a thread to wait for a condition and atomically release the associated mutex that it needs to hold to check the condition. The thread waits for another thread to make the condition true and that thread's resulting call to signal and wakeup the waiting thread.

The `cond_wait()` function atomically releases the mutex pointed to by *mp* and causes the calling thread to block on the condition variable pointed to by *cvp*. The blocked thread may be awakened by `cond_signal()`, `cond_broadcast()`, or when interrupted by delivery of a UNIX signal or a `fork()`.

The `cond_wait()`, `cond_timedwait()`, and `cond_reltimedwait()` functions always return with the mutex locked and owned by the calling thread even when returning an error, except when the mutex is of `USYNC_PROCESS_ROBUST` type and has been left irrecoverable by the mutex's last owner. The `cond_wait()`, `cond_timedwait()`, and `cond_reltimedwait()` functions return the appropriate error value if they fail to internally reacquire the mutex.

Condition Signaling

A condition signal allows a thread to unblock the next thread waiting on the condition variable, whereas, a condition broadcast allows a thread to unblock all threads waiting on the condition variable.

The `cond_signal()` function unblocks one thread that is blocked on the condition variable pointed to by *cvp*.

	<p>The <code>cond_broadcast()</code> function unblocks all threads that are blocked on the condition variable pointed to by <i>cvp</i>.</p> <p>If no threads are blocked on the condition variable, then <code>cond_signal()</code> and <code>cond_broadcast()</code> have no effect.</p> <p>Both functions should be called under the protection of the same mutex that is used with the condition variable being signaled. Otherwise, the condition variable may be signaled between the test of the associated condition and blocking in <code>cond_wait()</code>. This can cause an infinite wait.</p>										
Destroy	<p>The condition destroy functions destroy any state, but not the space, associated with the condition variable.</p> <p>The <code>cond_destroy()</code> function destroys any state associated with the condition variable pointed to by <i>cvp</i>. The space for storing the condition variable is not freed.</p>										
RETURN VALUES	<p>Upon successful completion, these functions return 0. Otherwise, a non-zero value is returned to indicate the error.</p>										
ERRORS	<p>These functions may fail if:</p> <table border="0"> <tr> <td style="vertical-align: top; padding-right: 20px;">EFAULT</td> <td>The <i>cond</i>, <i>attr</i>, <i>cvp</i>, <i>arg</i>, <i>abstime</i>, or <i>mutex</i> argument points to an illegal address.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">EINVAL</td> <td>Invalid argument. For <code>cond_init()</code>, <i>type</i> is not a recognized type. For <code>cond_timedwait()</code>, the specified number of seconds, <i>abstime</i>, is greater than <i>current_time</i> + 100,000,000, where <i>current_time</i> is the current time, or the number of nanoseconds is greater than or equal to 1,000,000,000.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">ENOTRECOVERABLE</td> <td>The mutex pointed to by <i>mp</i> is protecting the state that has been left irrecoverable when the mutex's last owner was not able to clean up the state. The mutex has not been acquired. See the description of <code>mutex_lock()</code> function on the <code>mutex_init(3THR)</code> manual page.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">EOWNERDEAD</td> <td>The last owner of the mutex pointed to by <i>mp</i> died while holding the mutex. The mutex has been acquired. See the description of <code>mutex_lock()</code> function on the <code>mutex_init(3THR)</code> manual page.</td> </tr> </table> <p>The <code>cond_timedwait()</code> and <code>cond_reltimedwait()</code> functions will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top; padding-right: 20px;">ETIME</td> <td>The time specified by <i>abstime</i> or <i>reltime</i> has passed.</td> </tr> </table>	EFAULT	The <i>cond</i> , <i>attr</i> , <i>cvp</i> , <i>arg</i> , <i>abstime</i> , or <i>mutex</i> argument points to an illegal address.	EINVAL	Invalid argument. For <code>cond_init()</code> , <i>type</i> is not a recognized type. For <code>cond_timedwait()</code> , the specified number of seconds, <i>abstime</i> , is greater than <i>current_time</i> + 100,000,000, where <i>current_time</i> is the current time, or the number of nanoseconds is greater than or equal to 1,000,000,000.	ENOTRECOVERABLE	The mutex pointed to by <i>mp</i> is protecting the state that has been left irrecoverable when the mutex's last owner was not able to clean up the state. The mutex has not been acquired. See the description of <code>mutex_lock()</code> function on the <code>mutex_init(3THR)</code> manual page.	EOWNERDEAD	The last owner of the mutex pointed to by <i>mp</i> died while holding the mutex. The mutex has been acquired. See the description of <code>mutex_lock()</code> function on the <code>mutex_init(3THR)</code> manual page.	ETIME	The time specified by <i>abstime</i> or <i>reltime</i> has passed.
EFAULT	The <i>cond</i> , <i>attr</i> , <i>cvp</i> , <i>arg</i> , <i>abstime</i> , or <i>mutex</i> argument points to an illegal address.										
EINVAL	Invalid argument. For <code>cond_init()</code> , <i>type</i> is not a recognized type. For <code>cond_timedwait()</code> , the specified number of seconds, <i>abstime</i> , is greater than <i>current_time</i> + 100,000,000, where <i>current_time</i> is the current time, or the number of nanoseconds is greater than or equal to 1,000,000,000.										
ENOTRECOVERABLE	The mutex pointed to by <i>mp</i> is protecting the state that has been left irrecoverable when the mutex's last owner was not able to clean up the state. The mutex has not been acquired. See the description of <code>mutex_lock()</code> function on the <code>mutex_init(3THR)</code> manual page.										
EOWNERDEAD	The last owner of the mutex pointed to by <i>mp</i> died while holding the mutex. The mutex has been acquired. See the description of <code>mutex_lock()</code> function on the <code>mutex_init(3THR)</code> manual page.										
ETIME	The time specified by <i>abstime</i> or <i>reltime</i> has passed.										
EXAMPLES	<p>EXAMPLE 1 Use <code>cond_wait()</code> in a loop to test some condition.</p> <p>The <code>cond_wait()</code> function is normally used in a loop testing some condition, as follows:</p>										

cond_init(3THR)

EXAMPLE 1 Use `cond_wait()` in a loop to test some condition. (Continued)

```
(void) mutex_lock(mp);
while (cond == FALSE) {
    (void) cond_wait(cvp, mp);
}
(void) mutex_unlock(mp);
```

EXAMPLE 2 Use `cond_timedwait()` in a loop to test some condition.

The `cond_timedwait()` function is normally used in a loop testing some condition. It uses an absolute timeout value as follows:

```
timestruc_t to;
...
(void) mutex_lock(mp);
to.tv_sec = time(NULL) + TIMEOUT;
to.tv_nsec = 0;
while (cond == FALSE) {
    err = cond_timedwait(cvp, mp, &to);
    if (err == ETIME) {
        /* timeout, do something */
        break;
    }
}
(void) mutex_unlock(mp);
```

EXAMPLE 3 Use `cond_reltimedwait()` in a loop to test some condition.

The `cond_reltimedwait()` function is normally used in a loop testing in some condition. It uses a relative timeout value as follows:

```
timestruc_t to;
...
(void) mutex_lock(mp);
while (cond == FALSE) {
    to.tv_sec = TIMEOUT;
    to.tv_nsec = 0;
    err = cond_reltimedwait(cvp, mp, &to);
    if (err == ETIME) {
        /* timeout, do something */
        break;
    }
}
(void) mutex_unlock(mp);
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

cond_init(3THR)

SEE ALSO fork(2), mmap(2), setitimer(2), shmop(2), condition(3THR), mutex(3THR), mutex_init(3THR), signal(3C), attributes(5), standards(5)

NOTES The only policy currently supported is SCHED_OTHER. In Solaris, under the SCHED_OTHER policy, there is no established order in which threads are unblocked.

If more than one thread is blocked on a condition variable, the order in which threads are unblocked is determined by the scheduling policy. When each thread, unblocked as a result of a cond_signal() or cond_broadcast(), returns from its call to cond_wait() or cond_timedwait(), the thread owns the mutex with which it called cond_wait(), cond_timedwait(), or cond_reltimedwait(). The thread(s) that are unblocked compete for the mutex according to the scheduling policy and as if each had called mutex_lock(3THR).

When cond_wait() returns the value of the condition is indeterminate and must be reevaluated.

The cond_timedwait() and cond_reltimedwait() functions are similar to cond_wait(), except that the calling thread will not wait for the condition to become true past the absolute time specified by *abstime* or the relative time specified by *reltime*. Note that cond_timedwait() or cond_reltimedwait() might continue to block as it tries to reacquire the mutex pointed to by *mp*, which may be locked by another thread. If either cond_timedwait() or cond_reltimedwait() returns because of a timeout, it returns the error value ETIME.

condition(3THR)

NAME	condition – concepts related to condition variables
DESCRIPTION	<p>Occasionally, a thread running within a mutex needs to wait for an event, in which case it blocks or sleeps. When a thread is waiting for another thread to communicate its disposition, it uses a condition variable in conjunction with a mutex. Although a mutex is exclusive and the code it protects is sharable (at certain moments), condition variables enable the synchronization of differing events that share a mutex, but not necessarily data. Several condition variables may be used by threads to signal each other when a task is complete, which then allows the next waiting thread to take ownership of the mutex.</p> <p>A condition variable enables threads to atomically block and test the condition under the protection of a mutual exclusion lock (mutex) until the condition is satisfied. If the condition is false, a thread blocks on a condition variable and atomically releases the mutex that is waiting for the condition to change. If another thread changes the condition, it may wake up waiting threads by signaling the associated condition variable. The waiting threads, upon awakening, reacquire the mutex and re-evaluate the condition.</p>
Initialize	<p>Condition variables and mutexes should be global. Condition variables that are allocated in writable memory can synchronize threads among processes if they are shared by the cooperating processes (see <code>mmap(2)</code>) and are initialized for this purpose.</p> <p>The scope of a condition variable is either intra-process or inter-process. This is dependent upon whether the argument is passed implicitly or explicitly to the initialization of that condition variable. A condition variable does not need to be explicitly initialized. A condition variable is initialized with all zeros, by default, and its scope is set to within the calling process. For inter-process synchronization, a condition variable must be initialized once, and only once, before use.</p> <p>A condition variable must not be simultaneously initialized by multiple threads or re-initialized while in use by other threads.</p> <p>Condition variables attributes may be set to the default or customized at initialization. POSIX threads even allow the default values to be customized. Establishing these attributes varies depending upon whether POSIX or Solaris threads are used. Similar to the distinctions between POSIX and Solaris thread creation, POSIX condition variables implement the default, intra-process, unless an attribute object is modified for inter-process prior to the initialization of the condition variable. Solaris condition variables also implement as the default, intra-process; however, they set this attribute according to the argument, <i>type</i>, passed to their initialization function.</p>
Condition Wait	<p>The condition wait interface allows a thread to wait for a condition and atomically release the associated mutex that it needs to hold to check the condition. The thread waits for another thread to make the condition true and that thread's resulting call to signal and wakeup the waiting thread.</p>

condition(3THR)

Condition Signaling A condition signal allows a thread to unblock the next thread waiting on the condition variable, whereas, a condition broadcast allows a thread to unblock all threads waiting on the condition variable.

Destroy The condition destroy functions destroy any state, but not the space, associated with the condition variable.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `fork(2)`, `mmap(2)`, `setitimer(2)`, `shmop(2)`, `cond_init(3THR)`, `cond_wait(3THR)`, `cond_timedwait(3THR)`, `cond_signal(3THR)`, `cond_broadcast(3THR)`, `cond_destroy(3THR)`, `mutex(3THR)`, `pthread_condattr_init(3THR)`, `pthread_cond_init(3THR)`, `pthread_cond_wait(3THR)`, `pthread_cond_timedwait(3THR)`, `pthread_cond_signal(3THR)`, `pthread_cond_broadcast(3THR)`, `pthread_cond_destroy(3THR)`, `signal(3C)`, `attributes(5)`, `standards(5)`

NOTES If more than one thread is blocked on a condition variable, the order in which threads are unblocked is determined by the scheduling policy.

`USYNC_THREAD` does not support multiple mappings to the same logical synch object. If you need to `mmap()` a synch object to different locations within the same address space, then the synch object should be initialized as a shared object `USYNC_PROCESS` for Solaris, and `PTHREAD_PROCESS_PRIVATE` for POSIX.

door_bind(3DOOR)

NAME	door_bind, door_unbind – bind or unbind the current thread with the door server pool						
SYNOPSIS	<pre>cc -mt [flag ...] file ... -ldoor [library ...] #include <door.h> int door_bind(int did); int door_unbind();</pre>						
DESCRIPTION	<p>The <code>door_bind()</code> function associates the current thread with a door server pool. A door server pool is a private pool of server threads that is available to serve door invocations associated with the door <i>did</i>.</p> <p>The <code>door_unbind()</code> function breaks the association of <code>door_bind()</code> by removing any private door pool binding that is associated with the current thread.</p> <p>Normally, door server threads are placed in a global pool of available threads that invocations on any door can use to dispatch a door invocation. A door that has been created with <code>DOOR_PRIVATE</code> only uses server threads that have been associated with the door by <code>door_bind()</code>. It is therefore necessary to bind at least one server thread to doors created with <code>DOOR_PRIVATE</code>.</p> <p>The server thread create function, <code>door_server_create()</code>, is initially called by the system during a <code>door_create()</code> operation. See <code>door_server_create(3DOOR)</code> and <code>door_create(3DOOR)</code>.</p> <p>The current thread is added to the private pool of server threads associated with a door during the next <code>door_return()</code> (that has been issued by the current thread after an associated <code>door_bind()</code>). See <code>door_return(3DOOR)</code>. A server thread performing a <code>door_bind()</code> on a door that is already bound to a different door performs an implicit <code>door_unbind()</code> of the previous door.</p> <p>If a process containing threads that have been bound to a door calls <code>fork(2)</code>, the threads in the child process will be bound to an invalid door, and any calls to <code>door_return(3DOOR)</code> will result in an error.</p>						
RETURN VALUES	Upon successful completion, a 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.						
ERRORS	The <code>door_bind()</code> and <code>door_unbind()</code> functions fail if:						
	<table><tr><td>EBADF</td><td>The <i>did</i> argument is not a valid door.</td></tr><tr><td>EBADF</td><td>The <code>door_unbind()</code> function was called by a thread that is currently not bound.</td></tr><tr><td>EINVAL</td><td><i>did</i> was not created with the <code>DOOR_PRIVATE</code> attribute.</td></tr></table>	EBADF	The <i>did</i> argument is not a valid door.	EBADF	The <code>door_unbind()</code> function was called by a thread that is currently not bound.	EINVAL	<i>did</i> was not created with the <code>DOOR_PRIVATE</code> attribute.
EBADF	The <i>did</i> argument is not a valid door.						
EBADF	The <code>door_unbind()</code> function was called by a thread that is currently not bound.						
EINVAL	<i>did</i> was not created with the <code>DOOR_PRIVATE</code> attribute.						

EXAMPLES **EXAMPLE 1** Use `door_bind()` to create private server pools for two doors.

The following example shows the use of `door_bind()` to create private server pools for two doors, `d1` and `d2`. Function `my_create()` is called when a new server thread is needed; it creates a thread running function, `my_server_create()`, which binds itself to one of the two doors.

```
#include <door.h>
#include <thread.h>
#include <pthread.h>
thread_key_t door_key;
int d1 = -1;
int d2 = -1;
cond_t cv;      /* statically initialized to zero */
mutex_t lock;   /* statically initialized to zero */

extern foo(); extern bar();

static void *
my_server_create(void *arg)
{
    /* wait for d1 & d2 to be initialized */
    mutex_lock(&lock);
    while (d1 == -1 || d2 == -1)
        cond_wait(&cv, &lock);
    mutex_unlock(&lock);

    if (arg == (void *)foo){
        /* bind thread with pool associated with d1 */
        thr_setspecific(door_key, (void *)foo);
        if (door_bind(d1) < 0) {
            perror("door_bind"); exit (-1);
        }
    } else if (arg == (void *)bar) {
        /* bind thread with pool associated with d2 */
        thr_setspecific(door_key, (void *)bar);
        if (door_bind(d2) < 0) {
            /* bind thread to d2 thread pool */
            perror("door_bind"); exit (-1);
        }
    }
    pthread_setcancelstate(POSIX_CANCEL_DISABLE, NULL);
    door_return(NULL, 0, NULL, 0); /* Wait for door invocation */
}

static void
my_create(door_info_t *dip)
{
    /* Pass the door identity information to create function */
    thr_create(NULL, 0, my_server_create, (void *)dip->di_proc,
              THR_BOUND | THR_DETACHED, NULL);
}

main( )
{
    (void)door_server_create(my_create);
    mutex_lock(&lock);
}
```

door_bind(3DOOR)

EXAMPLE 1 Use `door_bind()` to create private server pools for two doors. *(Continued)*

```
d1 = door_create(foo, NULL, DOOR_PRIVATE); /* Private pool */
d2 = door_create(bar, NULL, DOOR_PRIVATE); /* Private pool */
cond_signal(&cv);
mutex_unlock(&lock);
while (1)
    pause( );
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsu
Stability	Evolving
MT-Level	Safe

SEE ALSO `fork(2)`, `door_create(3DOOR)`, `door_return(3DOOR)`,
`door_server_create(3DOOR)`, `attributes(5)`

NAME	door_call – invoke the function associated with a door descriptor
SYNOPSIS	<pre>cc [flag ...] file ... -ldoor [library ...] #include <door.h> typedef struct { char *data_ptr; /* Argument/result buf ptr*/ size_t data_size; /* Argument/result buf size */ door_desc_t *desc_ptr; /* Argument/result descriptors */ uint_t desc_num; /* Argument/result num desc */ char *rbuf; /* Result buffer */ size_t rsize; /* Result buffer size */ } door_arg_t; int door_call(int d, door_arg_t *params);</pre>
DESCRIPTION	<p>The <code>door_call()</code> function invokes the function associated with the door descriptor <i>d</i>, and passes the arguments (if any) specified in <i>params</i>. All of the <i>params</i> members are treated as in/out parameters during a door invocation and may be updated upon returning from a door call. Passing NULL for <i>params</i> indicates there are no arguments to be passed and no results expected.</p> <p>Arguments are specified using the <code>data_ptr</code> and <code>desc_ptr</code> members of <i>params</i>. The size of the argument data in bytes is passed in <code>data_size</code> and the number of argument descriptors is passed in <code>desc_num</code>.</p> <p>Results from the door invocation are placed in the buffer, <code>rbuf</code>. See <code>door_return(3DOOR)</code>. The <code>data_ptr</code> and <code>desc_ptr</code> members of <i>params</i> are updated to reflect the location of the results within the <code>rbuf</code> buffer. The size of the data results and number of descriptors returned are updated in the <code>data_size</code> and <code>desc_num</code> members. It is acceptable to use the same buffer for input argument data and results, so <code>door_call()</code> may be called with <code>data_ptr</code> and <code>desc_ptr</code> pointing to the buffer <code>rbuf</code>.</p> <p>If the results of a door invocation exceed the size of the buffer specified by <code>rsize</code>, the system automatically allocates a new buffer in the caller's address space and updates the <code>rbuf</code> and <code>rsize</code> members to reflect this location. In this case, the caller is responsible for reclaiming this area using <code>munmap(rbuf, rsize)</code> when the buffer is no longer required. See <code>munmap(2)</code>.</p> <p>Descriptors passed in a <code>door_desc_t</code> structure are identified by the <code>d_attributes</code> member. The client marks the <code>d_attributes</code> member with the type of object being passed by logically OR-ing the value of object type. Currently, the only object type that can be passed or returned is a file descriptor, denoted by the <code>DOOR_DESCRIPTOR</code> attribute. Additionally, the <code>DOOR_RELEASE</code> attribute can be set, causing the descriptor to be closed in the caller's address space after it is passed to the target. The descriptor will be closed even if <code>door_call()</code> returns an error, unless that error is <code>EFAULT</code> or <code>EBADF</code>.</p> <p>The <code>door_desc_t</code> structure includes the following members:</p>

door_call(3DOOR)

```
typedef struct {
    door_attr_t d_attributes; /* Describes the parameter */
    union {
        struct {
            int d_descriptor; /* Descriptor */
            door_id_t d_id; /* Unique door id */
        } d_desc;
    } d_data;
} door_desc_t;
```

When file descriptors are passed or returned, a new descriptor is created in the target address space and the `d_descriptor` member in the target argument is updated to reflect the new descriptor. In addition, the system passes a system-wide unique number associated with each door in the `door_id` member and marks the `d_attributes` member with other attributes associated with a door including the following:

DOOR_LOCAL	The door received was created by this process using <code>door_create()</code> . See <code>door_create(3DOOR)</code> .
DOOR_PRIVATE	The door received has a private pool of server threads associated with the door.
DOOR_UNREF	The door received is expecting an unreferenced notification.
DOOR_UNREF_MULTI	Similar to <code>DOOR_UNREF</code> , except multiple unreferenced notifications may be delivered for the same door.
DOOR_REVOKED	The door received has been revoked by the server.

The `door_call()` function is not a restartable system call. It returns `EINTR` if a signal was caught and handled by this thread. If the door invocation is not idempotent the caller should mask any signals that may be generated during a `door_call()` operation. If the client aborts in the middle of a `door_call()`, the server thread is notified using the POSIX (see `standards(5)`) thread cancellation mechanism. See `cancellation(3THR)`.

The descriptor returned from `door_create()` is marked as close on exec (`FD_CLOEXEC`). Information about a door is available for all clients of a door using `door_info()`. Applications concerned with security should not place secure information in door data that is accessible by `door_info()`. In particular, secure data should not be stored in the data item `cookie`. See `door_info(3DOOR)`.

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

ERRORS The `door_call()` function will fail if:

EBADF	Invalid door descriptor was passed.
EINVAL	Bad arguments were passed.

door_call(3DOOR)

EFAULT Argument pointers pointed outside the allocated address space.
E2BIG Arguments were too big for server thread stack.
EOVERFLOW System could not create overflow area in caller for results.
EAGAIN Server was out of available resources.
EINTR Signal was caught in the client during the invocation.
EMFILE The client or server has too many open descriptors.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsu
Stability	Evolving
MT-Level	Safe

SEE ALSO `munmap(2)`, `cancellation(3THR)`, `door_create(3DOOR)`, `door_info(3DOOR)`, `door_return(3DOOR)`, `attributes(5)`, `standards(5)`

door_create(3DOOR)

NAME	door_create – create a door descriptor						
SYNOPSIS	<pre>cc -mt [flag ...] file ... -ldoor [library ...] #include <door.h> int door_create(void (*server_procedure) (void *cookie, char *argp, size_t arg_size, door_desc_t *dp, uint_t n_desc), void *cookie, uint_t attributes);</pre>						
DESCRIPTION	<p>The <code>door_create()</code> function creates a door descriptor that describes the procedure specified by the function <code>server_procedure</code>. The data item, <code>cookie</code>, is associated with the door descriptor, and is passed as an argument to the invoked function <code>server_procedure</code> during <code>door_call(3DOOR)</code> invocations. Other arguments passed to <code>server_procedure</code> from an associated <code>door_call()</code> are placed on the stack and include <code>argp</code> and <code>dp</code>. The <code>argp</code> argument points to <code>arg_size</code> bytes of data and the <code>dp</code> argument points to <code>n_desc</code> <code>door_desc_t</code> structures. The <code>attributes</code> argument specifies attributes associated with the newly created door. Valid values for <code>attributes</code> are constructed by OR-ing one or more of the following values:</p> <table><tr><td>DOOR_UNREF</td><td>Delivers a special invocation on the door when the number of descriptors that refer to this door drops to one. In order to trigger this condition, more than one descriptor must have referred to this door at some time. <code>DOOR_UNREF_DATA</code> designates an unreferenced invocation, as the <code>argp</code> argument passed to <code>server_procedure</code>. In the case of an unreferenced invocation, the values for <code>arg_size</code>, <code>dp</code> and <code>n_desc</code> are 0. Only one unreferenced invocation is delivered on behalf of a door.</td></tr><tr><td>DOOR_UNREF_MULTI</td><td>Similar to <code>DOOR_UNREF</code>, except multiple unreferenced invocations can be delivered on the same door if the number of descriptors referring to the door drops to one more than once. Since an additional reference may have been passed by the time an unreferenced invocation arrives, the <code>DOOR_IS_UNREF</code> attribute returned by the <code>door_info(3DOOR)</code> call can be used to determine if the door is still unreferenced.</td></tr><tr><td>DOOR_PRIVATE</td><td>Maintains a separate pool of server threads on behalf of the door. Server threads are associated with a door's private server pool using <code>door_bind(3DOOR)</code>.</td></tr></table> <p>The descriptor returned from <code>door_create()</code> will be marked as close on exec (<code>FD_CLOEXEC</code>). Information about a door is available for all clients of a door using <code>door_info(3DOOR)</code>. Applications concerned with security should not place secure information in door data that is accessible by <code>door_info()</code>. In particular, secure data should not be stored in the data item <code>cookie</code>.</p>	DOOR_UNREF	Delivers a special invocation on the door when the number of descriptors that refer to this door drops to one. In order to trigger this condition, more than one descriptor must have referred to this door at some time. <code>DOOR_UNREF_DATA</code> designates an unreferenced invocation, as the <code>argp</code> argument passed to <code>server_procedure</code> . In the case of an unreferenced invocation, the values for <code>arg_size</code> , <code>dp</code> and <code>n_desc</code> are 0. Only one unreferenced invocation is delivered on behalf of a door.	DOOR_UNREF_MULTI	Similar to <code>DOOR_UNREF</code> , except multiple unreferenced invocations can be delivered on the same door if the number of descriptors referring to the door drops to one more than once. Since an additional reference may have been passed by the time an unreferenced invocation arrives, the <code>DOOR_IS_UNREF</code> attribute returned by the <code>door_info(3DOOR)</code> call can be used to determine if the door is still unreferenced.	DOOR_PRIVATE	Maintains a separate pool of server threads on behalf of the door. Server threads are associated with a door's private server pool using <code>door_bind(3DOOR)</code> .
DOOR_UNREF	Delivers a special invocation on the door when the number of descriptors that refer to this door drops to one. In order to trigger this condition, more than one descriptor must have referred to this door at some time. <code>DOOR_UNREF_DATA</code> designates an unreferenced invocation, as the <code>argp</code> argument passed to <code>server_procedure</code> . In the case of an unreferenced invocation, the values for <code>arg_size</code> , <code>dp</code> and <code>n_desc</code> are 0. Only one unreferenced invocation is delivered on behalf of a door.						
DOOR_UNREF_MULTI	Similar to <code>DOOR_UNREF</code> , except multiple unreferenced invocations can be delivered on the same door if the number of descriptors referring to the door drops to one more than once. Since an additional reference may have been passed by the time an unreferenced invocation arrives, the <code>DOOR_IS_UNREF</code> attribute returned by the <code>door_info(3DOOR)</code> call can be used to determine if the door is still unreferenced.						
DOOR_PRIVATE	Maintains a separate pool of server threads on behalf of the door. Server threads are associated with a door's private server pool using <code>door_bind(3DOOR)</code> .						

door_create(3DOOR)

By default, additional threads are created as needed to handle concurrent door_call(3DOOR) invocations. See door_server_create(3DOOR) for information on how to change this behavior.

RETURN VALUES Upon successful completion, door_create() returns a non-negative value. Otherwise, door_create returns -1 and sets errno to indicate the error.

ERRORS The door_create() function will fail if:

EINVAL Invalid attributes are passed.

EMFILE The process has too many open descriptors.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsu
Interface Stability	Evolving
MT-Level	Safe

SEE ALSO door_bind(3DOOR), door_call(3DOOR), door_info(3DOOR), door_revoke(3DOOR), door_server_create(3DOOR), fattach(3C), attributes(5)

door_cred(3DOOR)

NAME door_cred – return credential information associated with the client

SYNOPSIS

```
cc -mt [ flag ... ] file ... -ldoor [ library ... ]
#include <door.h>

int door_cred(door_cred_t *info);
```

DESCRIPTION The door_cred() function returns credential information associated with the client (if any) of the current door invocation.

The contents of the *info* argument include the following fields:

```
uid_t  dc_euid;      /* Effective uid of client */
gid_t  dc_egid;     /* Effective gid of client */
uid_t  dc_ruid;     /* Real uid of client */
gid_t  dc_rgid;     /* Real gid of client */
pid_t  dc_pid;      /* pid of client */
```

The credential information associated with the client refers to the information from the immediate caller; not necessarily from the first thread in a chain of door calls.

RETURN VALUES Upon successful completion, door_cred() returns 0. Otherwise, door_cred() returns -1 and sets errno to indicate the error.

ERRORS The door_cred() function will fail if:

EFAULT	The address of the <i>info</i> argument is invalid.
EINVAL	There is no associated door client.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsu
Stability	Evolving
MT-Level	Safe

SEE ALSO door_call(3DOOR), door_create(3DOOR), attributes(5)

NAME	door_info – return information associated with a door descriptor												
SYNOPSIS	<pre>cc [flag ...] file ... -ldoor [library ...] #include <door.h> int door_info(int d, struct door_info *info);</pre>												
DESCRIPTION	<p>The <code>door_info()</code> function returns information associated with a door descriptor. It obtains information about the door descriptor <code>d</code> and places the information that is relevant to the door in the structure pointed to by the <code>info</code> argument.</p> <p>The <code>door_info</code> structure pointed to by the <code>info</code> argument contains the following members:</p> <pre>pid_t di_target; /* door server pid */ door_ptr_t di_proc; /* server function */ door_ptr_t di_data; /* data cookie for invocation */ door_attr_t di_attributes; /* door attributes */ door_id_t di_uniquifier; /* unique id among all doors */</pre> <p>The <code>di_target</code> member is the process ID of the door server, or <code>-1</code> if the door server process has exited.</p> <p>The values for <code>di_attributes</code> may be composed of the following:</p> <table border="0"> <tr> <td style="padding-right: 20px;">DOOR_LOCAL</td> <td>The door descriptor refers to a service procedure in this process.</td> </tr> <tr> <td>DOOR_UNREF</td> <td>The door has requested notification when all but the last reference has gone away.</td> </tr> <tr> <td>DOOR_UNREF_MULTI</td> <td>Similar to <code>DOOR_UNREF</code>, except multiple unreferenced notifications may be delivered for this door.</td> </tr> <tr> <td>DOOR_IS_UNREF</td> <td>There is currently only one descriptor referring to the door.</td> </tr> <tr> <td>DOOR_REVOKED</td> <td>The door descriptor refers to a door that has been revoked.</td> </tr> <tr> <td>DOOR_PRIVATE</td> <td>The door has a separate pool of server threads associated with it.</td> </tr> </table> <p>The <code>di_proc</code> and <code>di_data</code> members are returned as <code>door_ptr_t</code> objects rather than <code>void *</code> pointers to allow clients and servers to interoperate in environments where the pointer sizes may vary in size (for example, 32-bit clients and 64-bit servers). Each door has a system-wide unique number associated with it that is set when the door is created by <code>door_create()</code>. This number is returned in <code>di_uniquifier</code>.</p>	DOOR_LOCAL	The door descriptor refers to a service procedure in this process.	DOOR_UNREF	The door has requested notification when all but the last reference has gone away.	DOOR_UNREF_MULTI	Similar to <code>DOOR_UNREF</code> , except multiple unreferenced notifications may be delivered for this door.	DOOR_IS_UNREF	There is currently only one descriptor referring to the door.	DOOR_REVOKED	The door descriptor refers to a door that has been revoked.	DOOR_PRIVATE	The door has a separate pool of server threads associated with it.
DOOR_LOCAL	The door descriptor refers to a service procedure in this process.												
DOOR_UNREF	The door has requested notification when all but the last reference has gone away.												
DOOR_UNREF_MULTI	Similar to <code>DOOR_UNREF</code> , except multiple unreferenced notifications may be delivered for this door.												
DOOR_IS_UNREF	There is currently only one descriptor referring to the door.												
DOOR_REVOKED	The door descriptor refers to a door that has been revoked.												
DOOR_PRIVATE	The door has a separate pool of server threads associated with it.												
RETURN VALUES	Upon successful completion, <code>0</code> is returned. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.												
ERRORS	The <code>door_info()</code> function will fail if:												

door_info(3DOOR)

EFAULT The address of argument *info* is an invalid address.
EBADF *d* is not a door descriptor.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsu
Stability	Evolving
MT-Level	Safe

SEE ALSO `door_bind(3DOOR)`, `door_create(3DOOR)`, `door_server_create(3DOOR)`

NAME	door_return – return from a door invocation										
SYNOPSIS	<pre>cc -mt [flag ...] file ... -ldoor [library ...] #include <door.h> int door_return(char *data_ptr, size_t data_size, door_desc_t *desc_ptr, uint_t num_desc);</pre>										
DESCRIPTION	<p>The <code>door_return()</code> function returns from a door invocation. It returns control to the thread that issued the associated <code>door_call()</code> and blocks waiting for the next door invocation. See <code>door_call(3DOOR)</code>. Results, if any, from the door invocation are passed back to the client in the buffers pointed to by <code>data_ptr</code> and <code>desc_ptr</code>. If there is not a client associated with the <code>door_return()</code>, the calling thread discards the results, releases any passed descriptors with the <code>DOOR_RELEASE</code> attribute, and blocks waiting for the next door invocation.</p>										
RETURN VALUES	<p>Upon successful completion, <code>door_return()</code> does not return to the calling process. Otherwise, <code>door_return()</code> returns <code>-1</code> to the calling process and sets <code>errno</code> to indicate the error.</p>										
ERRORS	<p>The <code>door_return()</code> function fails and returns to the calling process if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>E2BIG</code></td> <td>Arguments were too big for client.</td> </tr> <tr> <td><code>EFAULT</code></td> <td>The address of <code>data_ptr</code> or <code>desc_ptr</code> is invalid.</td> </tr> <tr> <td><code>EINVAL</code></td> <td>Invalid <code>door_return()</code> arguments were passed or a thread is bound to a door that no longer exists.</td> </tr> <tr> <td><code>EMFILE</code></td> <td>The client has too many open descriptors.</td> </tr> </table>	<code>E2BIG</code>	Arguments were too big for client.	<code>EFAULT</code>	The address of <code>data_ptr</code> or <code>desc_ptr</code> is invalid.	<code>EINVAL</code>	Invalid <code>door_return()</code> arguments were passed or a thread is bound to a door that no longer exists.	<code>EMFILE</code>	The client has too many open descriptors.		
<code>E2BIG</code>	Arguments were too big for client.										
<code>EFAULT</code>	The address of <code>data_ptr</code> or <code>desc_ptr</code> is invalid.										
<code>EINVAL</code>	Invalid <code>door_return()</code> arguments were passed or a thread is bound to a door that no longer exists.										
<code>EMFILE</code>	The client has too many open descriptors.										
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>all</td> </tr> <tr> <td>Availability</td> <td>SUNWcsu</td> </tr> <tr> <td>Stability</td> <td>Evolving</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	all	Availability	SUNWcsu	Stability	Evolving	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE										
Architecture	all										
Availability	SUNWcsu										
Stability	Evolving										
MT-Level	Safe										
SEE ALSO	<code>door_call(3DOOR)</code>										

door_revoke(3DOOR)

NAME | door_revoke – revoke access to a door descriptor

SYNOPSIS |

```
cc -mt [ flag ... ] file ... -ldoor [ library ... ]
#include <door.h>

int door_revoke(int d);
```

DESCRIPTION | The door_revoke() function revokes access to a door descriptor. Door descriptors are created with door_create(3DOOR). The door_revoke() function performs an implicit call to close(2), marking the door descriptor *d* as invalid.

A door descriptor can only be revoked by the process that created it. Door invocations that are in progress during a door_revoke() invocation are allowed to complete normally.

RETURN VALUES | Upon successful completion, door_revoke() returns 0. Otherwise, door_revoke() returns -1 and sets errno to indicate the error.

ERRORS | The door_revoke() function will fail if:

EBADF	An invalid door descriptor was passed.
EPERM	The door descriptor was not created by this process (with door_create(3DOOR)).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsu
Stability	Evolving
MT-Level	Safe

SEE ALSO | close(2), door_create(3DOOR), attributes(5)

NAME	door_server_create – specify an alternative door server thread creation function
SYNOPSIS	<pre>cc -mt [flag ...] file ... -ldoor [library ...] #include <door.h> void (*) () door_server_create(void (*create_proc)(door_info_t*));</pre>
DESCRIPTION	<p>Normally, the doors library creates new door server threads in response to incoming concurrent door invocations automatically. There is no pre-defined upper limit on the number of server threads that the system creates in response to incoming invocations (1 server thread for each active door invocation). These threads are created with the default thread stack size and POSIX (see standards(5)) threads cancellation disabled. The created threads also have the THR_BOUND THR_DETACHED attributes for Solaris threads and the PTHREAD_SCOPE_SYSTEM PTHREAD_CREATE_DETACHED attributes for POSIX threads. The signal disposition, and scheduling class of the newly created thread are inherited from the calling thread (initially from the thread calling door_create(), and subsequently from the current active door server thread).</p> <p>The door_server_create() function allows control over the creation of server threads needed for door invocations. The procedure create_proc is called every time the available server thread pool is depleted. In the case of private server pools associated with a door (see the DOOR_PRIVATE attribute in door_create()), information on which pool is depleted is passed to the create function in the form of a door_info_t structure. The di_proc and di_data members of the door_info_t structure can be used as a door identifier associated with the depleted pool. The create_proc procedure may limit the number of server threads created and may also create server threads with appropriate attributes (stack size, thread-specific data, POSIX thread cancellation, signal mask, scheduling attributes, and so forth) for use with door invocations.</p> <p>The specified server creation function should create user level threads using thr_create() with the THR_BOUND flag, or in the case of POSIX threads, pthread_create() with the PTHREAD_SCOPE_SYSTEM attribute. The server threads make themselves available for incoming door invocations on this process by issuing a door_return(NULL, 0, NULL, 0). In this case, the door_return() arguments are ignored. See door_return(3DOOR) and thr_create(3THR).</p> <p>The server threads created by default are enabled for POSIX thread cancellations which may lead to unexpected thread terminations while holding resources (such as locks) if the client aborts the associated door_call(). See door_call(3DOOR). Unless the server code is truly interested in notifications of client aborts during a door invocation and is prepared to handle such notifications using cancellation handlers, POSIX thread cancellation should be disabled for server threads using pthread_setcancelstate(PTHREAD_CANCEL_DISABLE, NULL).</p> <p>The create_proc procedure need not create any additional server threads if there is at least one server thread currently active in the process (perhaps handling another door invocation) or it may create as many as seen fit each time it is called. If there are no available server threads during an incoming door invocation, the associated</p>

door_server_create(3DOOR)

door_call() blocks until a server thread becomes available. The *create_proc* procedure must be MT-Safe.

RETURN VALUES Upon successful completion, door_server_create() returns a pointer to the previous server creation function. This function has no failure mode (it cannot fail).

EXAMPLES **EXAMPLE 1** Creating door server threads.

The following example creates door server threads with cancellation disabled and an 8k stack instead of the default stack size:

```
#include <door.h>
#include <pthread.h>
#include <thread.h>

void *
my_thread(void *arg)
{
    pthread_setcancelstate(PTHREAD_CANCEL_DISABLE, NULL);
    door_return(NULL, 0, NULL, 0);
}

void
my_create(door_info_t *dip)
{
    thr_create(NULL, 8192, my_thread, NULL,
              THR_BOUND | THR_DETACHED, NULL);
}

main( )
{
    (void) door_server_create(my_create);
    . . .
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsu
Stability	Evolving
MT-Level	Safe

SEE ALSO cancellation(3THR), door_bind(3DOOR), door_call(3DOOR), door_create(3DOOR), door_return(3DOOR), pthread_create(3THR), pthread_setcancelstate(3THR), thr_create(3THR), attributes(5), standards(5)

NAME	fdatasync – synchronize a file's data				
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <unistd.h> int fdatasync(int <i>fildev</i>);</pre>				
DESCRIPTION	<p>The <code>fdatasync()</code> function forces all currently queued I/O operations associated with the file indicated by file descriptor <i>fildev</i> to the synchronized I/O completion state.</p> <p>The functionality is as described for <code>fsync(3C)</code> (with the symbol <code>_XOPEN_REALTIME</code> defined), with the exception that all I/O operations are completed as defined for synchronised I/O data integrity completion.</p>				
RETURN VALUES	If successful, the <code>fdatasync()</code> function returns 0. Otherwise, the function returns <code>-1</code> and sets <code>errno</code> to indicate the error. If the <code>fdatasync()</code> function fails, outstanding I/O operations are not guaranteed to have been completed.				
ERRORS	<p>The <code>fdatasync()</code> function will fail if:</p> <p><code>EBADF</code> The <i>fildev</i> argument is not a valid file descriptor open for writing.</p> <p><code>EINVAL</code> The system does not support synchronized I/O for this file.</p> <p><code>ENOSYS</code> The function <code>fdatasync()</code> is not supported by the system.</p> <p>In the event that any of the queued I/O operations fail, <code>fdatasync()</code> returns the error conditions defined for <code>read(2)</code> and <code>write(2)</code>.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	<code>fcntl(2)</code> , <code>open(2)</code> , <code>read(2)</code> , <code>write(2)</code> , <code>fsync(3C)</code> , <code>aio_fsync(3RT)</code> , <code>attributes(5)</code> , <code>fcntl(3HEAD)</code>				

libthread_db(3THR)

NAME	libthread_db – library of interfaces for monitoring and manipulating threads-related aspects of multithreaded programs
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i> ... -lthread_db [<i>library</i> ...] #include <proc_service.h> #include <thread_db.h> void td_event_addset(td_thr_events_t *, td_thr_events_e n); void td_event_delset(td_thr_events_t *, td_thr_events_e n); void td_event_emptyset(td_thr_events_t *); void td_event_fillset(td_thr_events_t *); void td_eventisempty(td_thr_events_t *); void td_eventismember(td_thr_events_t *, td_thr_events_e n); td_err_e td_init(); void td_log(); td_err_e td_sync_get_info(const td_synchandle_t *sh_p, td_syncstats_t *ss_p); td_err_e td_sync_get_stats(const td_synchandle_t *sh_p, td_syncinfo_t *si_p); td_err_e td_sync_setstate(const td_synchandle_t *sh_p, int <i>value</i>); td_err_e td_sync_waiters(const td_synchandle_t *sh_p, td_thr_iter_f *cb, void *cb_data_p); td_err_e td_thr_clear_event(const td_thrhandle_t *th_p, td_thr_events_t *events); td_err_e td_ta_delete(td_thragent_t *ta_p); td_err_e td_ta_enable_stats(const td_thragent_t *ta_p, int <i>on_off</i>); td_err_e td_ta_event_addr(const td_thragent_t *ta_p, u_long event, td_notify_t *notify_p); td_err_e td_ta_event_getmsg(const td_thragent_t *ta_p, td_event_msg_t *msg); td_err_e td_ta_get_nthreads(const td_thragent_t *ta_p, int *nthread_p); td_err_e td_ta_get_ph(const td_thragent_t *ta_p, struct ps_prochandle **ph_pp); td_err_e td_ta_get_stats(const td_thragent_t *ta_p, td_ta_stats_t *tstats);</pre>

```

td_err_e td_ta_map_addr2sync(const td_thragent_t *ta_p, psaddr_t
    addr td_synchandle_t *sh_p);
td_err_e td_ta_map_id2thr(const td_thragent_t *ta_p, thread_t tid,
    td_thrhandle_t *th_p);
td_err_e td_ta_map_lwp2thr(const td_thragent_t *ta_p, lwpid_t
    lwpid, td_thrhandle_t *th_p);
td_err_e td_ta_new(struct ps_prochandle *ph_p, td_thragent_t
    **ta_pp);
td_err_e td_ta_reset_stats(const td_thragent_t *ta_p);
td_err_e td_ta_setconcurrency(const td_thragent_t *ta_p, int level);
td_err_e td_ta_sync_iter(const td_thragent_t *ta_p, td_sync_iter_f
    *cb, void *cbdata_p);
td_err_e td_ta_sync_tracking_enable(const td_thragent_t *ta_p, int
    on_off);
td_err_e td_ta_thr_iter(const td_thragent_t *ta_p, td_key_iter_f
    *cb, void *cbdata_p);
td_err_e td_ta_tsd_iter(const td_thragent_t *ta_p, td_key_iter_f
    *cb, void *cbdata_p);
td_err_e td_thr_clear_event(const td_thrhandle_t *th_p,
    td_thr_events_t *events);
td_err_e td_thr_dbresume(const td_thrhandle_t *th_p);
td_err_e td_thr_dbsuspend(const td_thrhandle_t *th_p);
td_err_e td_thr_event_enable(const td_thrhandle_t *th_p, int
    on_off);
td_err_e td_thr_event_getmsg(const td_thrhandle_t,
    td_event_msg_t *msg);
td_err_e td_thr_get_info(const td_thrhandle_t *th_p, td_thrinfo_t
    *ti_p);
td_err_e td_thr_getfpregs(const td_thrhandle_t *th_p,
    prfpregset_t *fpregset);
td_err_e td_thr_getgregs(const td_thrhandle_t *th_p, prgregset_t
    regset);
td_err_e td_thr_getxregs(const td_thrhandle_t *th_p, void *xregset);
td_err_e td_thr_getxregsize(const td_thrhandle_t *th_p, int
    *xregsize);

```

libthread_db(3THR)

```
td_err_e td_thr_lockowner(const td_thrhandle_t *th_p,  
    td_sync_iter_f *cb, void *cb_data_p);  
  
td_err_e td_thr_set_event(const td_thrhandle_t *th_p,  
    td_thr_events_t *events);  
  
td_err_e td_thr_setfpregs(const td_thrhandle_t *th_p,  
    prfpregset_t *fpregset);  
  
td_err_e td_thr_setgregs(const td_thrhandle_t *th_p, const  
    prgregset_t regset);  
  
td_err_e td_thr_setprio(const td_thrhandle_t *th_p, const int  
    new_prio);  
  
td_err_e td_thr_setsigpending(const td_thrhandle_t *th_p, const  
    uchar_t, ti_pending_flag, const sigset_t ti_pending);  
  
td_err_e td_thr_setxregs(const td_thrhandle_t *th_p, const void  
    *xregset);  
  
td_err_e td_thr_sigsetmask(const td_thrhandle_t *th_p, const  
    sigset_t ti_sigmask);  
  
td_err_e td_thr_sleepinfo(const td_thrhandle_t *th_p,  
    td_synchandle_t *sh_p);  
  
td_err_e td_thr_tsd(const td_thrhandle_t *th_p, const thread_key_t  
    key, void **data_pp);  
  
td_err_e td_thr_validate(const td_thrhandle_t *th_p);
```

DESCRIPTION

The `libthread_db` library provides support for monitoring and manipulating threads-related aspects of a multithreaded program. There are at least two processes involved, the controlling process and one or more target processes. The controlling process is the `libthread_db` client, which links with `libthread_db` and uses `libthread_db` to inspect or modify threads-related aspects of one or more target processes. The target processes must be multithreaded processes that use `libthread` or `libpthread`. The controlling process may or may not be multithreaded itself.

The most commonly anticipated use for `libthread_db` is that the controlling process will be a debugger for a multithreaded program, hence the "db" in `libthread_db`.

`libthread_db` is dependent on the internal implementation details of `libthread`. It is a "friend" of `libthread` in the C++ sense, which is precisely the "value added" by `libthread_db`. It encapsulates the knowledge of `libthread` internals that a debugger needs in order to manipulate the threads-related state of a target process.

To be able to inspect and manipulate target processes, `libthread_db` makes use of certain process control primitives that must be provided by the process using `libthread_db`. The imported interfaces are defined in `proc_service(3PROC)`. In other words, the controlling process is linked with `libthread_db`, and it calls

routines in `libthread_db`. `libthread_db` in turn calls certain routines that it expects the controlling process to provide. These process control primitives allow `libthread_db` to:

- Look up symbols in a target process.
- Stop and continue individual lightweight processes (LWPs) within a target process.
- Stop and continue an entire target process.
- Read and write memory and registers in a target process.

Initially, a controlling process obtains a handle for a target process. Through that handle it can then obtain handles for the component objects of the target process, its threads, its synchronization objects, and its thread-specific-data keys.

When `libthread_db` needs to return sets of handles to the controlling process, for example, when returning handles for all the threads in a target process, it uses an iterator function. An iterator function calls back a client-specified function once for each handle to be returned, passing one handle back on each call to the callback function. The calling function also passes another parameter to the iterator function, which the iterator function passes on to the callback function. This makes it easy to build a linked list of thread handles for a particular target process. The additional parameter is the head of the linked list, and the callback function simply inserts the current handle into the linked list.

Callback functions are expected to return an integer. Iteration terminates early if a callback function returns a non-zero value. Otherwise, iteration terminates when there are no more handles to pass back.

`libthread_db` relies on an "agent thread" in the target process for some of its operations. The "agent thread" is a system thread started when `libthread_db` attaches to a process through `td_ta_new(3THR)`. In the current implementation, a brief window exists after the agent thread has been started, but before it has completed its initialization, in which `libthread_db` routines that require the agent thread will fail, returning a `TD_NOCAPAB` error status. This is particularly troublesome if the target process was stopped when `td_ta_new()` was called, so that the agent thread cannot be initialized. To avoid this problem, the target process must be allowed to make some forward progress after `td_ta_new()` is called. This limitation will be removed in a future release.

FUNCTIONS

Name	Description
<code>td_event_addset()</code>	Macro that adds a specific event type to an event set.
<code>td_event_delset()</code>	Macro that deletes a specific event type from an event set.
<code>td_event_emptyset()</code>	Macro that sets argument to NULL event set.

libthread_db(3THR)

<code>td_event_fillset()</code>	Macro that sets argument to set of all events.
<code>td_eventisempty()</code>	Macro that tests whether an event set is the NULL set.
<code>td_eventismember()</code>	Macro that tests whether a specific event type is a member of an event set.
<code>td_init()</code>	Performs initialization for interfaces.
<code>td_log()</code>	Placeholder for future logging functionality.
<code>td_sync_get_info()</code>	Gets information for the synchronization object.
<code>td_sync_get_stats()</code>	Gets statistics for the synchronization object.
<code>td_sync_setstate()</code>	Sets the state of the synchronization object.
<code>td_sync_waiters()</code>	Iteration function used for return of synchronization object handles.
<code>td_ta_clear_event()</code>	Clears a set of event types in the process event mask.
<code>td_ta_delete()</code>	Deregisters target process and deallocates internal process handle.
<code>td_ta_enable_stats()</code>	Turns statistics gathering on or off for the target process.
<code>td_ta_event_addr()</code>	Returns event reporting address.
<code>td_ta_event_getmsg()</code>	Returns process event message.
<code>td_ta_get_nthreads()</code>	Gets the total number of threads in a process. .
<code>td_ta_get_ph()</code>	Returns corresponding external process handle.
<code>td_ta_get_stats()</code>	Gets statistics gathered for the target process.
<code>td_ta_map_addr2sync()</code>	Gets a synchronization object handles from a synchronization object's address.
<code>td_ta_map_id2thr()</code>	Returns a thread handle for the given thread id.
<code>td_ta_map_lwp2thr()</code>	Returns a thread handle for the given LWP id.

<code>td_ta_new()</code>	Registers target process and allocates internal process handle.
<code>td_ta_reset_stats()</code>	Resets all counters for statistics gathering for the target process.
<code>td_ta_setconcurrency()</code>	Sets concurrency level for target process.
<code>td_ta_set_event()</code>	Sets a set of event types in the process event mask.
<code>td_ta_sync_iter()</code>	Returns handles of synchronization objects associated with a process.
<code>td_ta_sync_tracking_enable()</code>	Enables or disables synchronization object tracking.
<code>td_ta_thr_iter()</code>	Returns handles for threads that are part of the target process.
<code>td_ta_tsd_iter()</code>	Returns the thread-specific data keys in use by the current process.
<code>td_thr_clear_event()</code>	Clears a set of event types in the threads event mask.
<code>td_thr_dbresume()</code>	Resumes thread.
<code>td_thr_dbsuspend()</code>	Suspends thread.
<code>td_thr_event_enable()</code>	Enables or disables event reporting.
<code>td_thr_event_getmsg()</code>	Returns a process event message.
<code>td_thr_get_info()</code>	Gets thread information and updates
<code>td_thr_getfpregs()</code>	Gets the floating point registers for the given thread.
<code>td_thr_getgregs()</code>	Gets the general registers for a given thread.
<code>td_thr_getxregs()</code>	Gets the extra registers for the given thread.
<code>td_thr_getxregsize()</code>	Gets the size of the extra register set for the given thread.
<code>td_thr_lockowner()</code>	Iterates over the set of locks owned by a thread. <code>struct</code> .
<code>td_thr_set_event()</code>	Sets a set of event types in the threads event mask.
<code>td_thr_setfpregs()</code>	Sets the floating point registers for the given thread. <code>ti_sigmask</code>
<code>td_thr_setgregs()</code>	Sets the general registers for a given thread.

libthread_db(3THR)

<code>td_thr_setprio()</code>	Sets the priority of a thread.
<code>td_thr_setsigpending()</code>	Changes a thread's pending signal state.
<code>td_thr_setxregs()</code>	Sets the extra registers for the given thread.
<code>td_thr_sigsetmask()</code>	Sets the signal mask of the thread.
<code>td_thr_sleepinfo()</code>	Returns the synchronization handle for the object on which a thread is blocked.
<code>td_thr_tsd()</code>	Gets a thread's thread-specific data.
<code>td_thr_validate()</code>	Tests a thread handle for validity.

FILES

`lthread_db`

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO

`libthread(3LIB)`, `libthread_db(3LIB)`, `proc_service(3PROC)`, `rtld_db(3EXT)`, `td_event_addset(3THR)`, `td_event_delset(3THR)`, `td_event_emptyset(3THR)`, `td_event_fillset(3THR)`, `td_eventisempty(3THR)`, `td_eventismember(3THR)`, `td_init(3THR)`, `td_log(3THR)`, `td_sync_get_info(3THR)`, `td_sync_get_stats(3THR)`, `td_sync_waiters(3THR)`, `td_ta_delete(3THR)`, `td_ta_enable_stats(3THR)`, `td_ta_event_addr(3THR)`, `td_ta_event_getmsg(3THR)`, `td_ta_get_nthreads(3THR)`, `td_ta_get_ph(3THR)`, `td_ta_get_stats(3THR)`, `td_ta_map_addr2sync(3THR)`, `td_ta_map_id2thr(3THR)`, `td_ta_map_lwp2thr(3THR)`, `td_ta_new(3THR)`, `td_ta_reset_stats(3THR)`, `td_ta_set_event(3THR)`, `td_ta_setconcurrency(3THR)`, `td_ta_sync_iter(3THR)`, `td_ta_sync_iter(3THR)`, `td_ta_sync_tracking_enable(3THR)`, `td_ta_tsd_iter(3THR)`, `td_thr_clear_event(3THR)`, `td_thr_dbresume(3THR)`, `td_thr_dbsuspend(3THR)`, `td_thr_event_enable(3THR)`, `td_thr_event_getmsg(3THR)`, `td_thr_get_info(3THR)`, `td_thr_getfpregs(3THR)`, `td_thr_getxregs(3THR)`, `td_thr_getxregsize(3THR)`, `td_thr_lockowner(3THR)`, `td_thr_set_event(3THR)`, `td_thr_setfpregs(3THR)`, `td_thr_setgregs(3THR)`, `td_thr_setprio(3THR)`, `td_thr_sigsetmask(3THR)`, `td_thr_setsigpending(3THR)`, `td_thr_setxregs(3THR)`, `td_thr_sleepinfo(3THR)`, `td_thr_tsd(3THR)`, `td_thr_validate(3THR)`, `thr_getspecific(3THR)`, `threads(3THR)`, `attributes(5)`

NAME	lio_listio – list directed I/O
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <aio.h> int lio_listio(int mode, struct aiocb * const list[], int nent, struct sigevent *sig);</pre>
DESCRIPTION	<p>The <code>lio_listio()</code> function allows the calling process, LWP, or thread, to initiate a list of I/O requests within a single function call.</p> <p>The <code>mode</code> argument takes one of the values <code>LIO_WAIT</code> or <code>LIO_NOWAIT</code> declared in <code><aio.h></code> and determines whether the function returns when the I/O operations have been completed, or as soon as the operations have been queued. If the <code>mode</code> argument is <code>LIO_WAIT</code>, the function waits until all I/O is complete and the <code>sig</code> argument is ignored.</p> <p>If the <code>mode</code> argument is <code>LIO_NOWAIT</code>, the function returns immediately, and asynchronous notification occurs, according to the <code>sig</code> argument, when all the I/O operations complete. If <code>sig</code> is <code>NULL</code>, or the <code>sigev_signo</code> member of the <code>sigevent</code> structure referenced by <code>sig</code> is zero, then no asynchronous notification occurs. If <code>sig</code> is not <code>NULL</code>, asynchronous notification occurs when all the requests in <code>list</code> have completed. If <code>sig->sigev_notify</code> is <code>SIGEV_NONE</code>, then no signal will be posted upon I/O completion, but the error status and the return status for the operation will be set appropriately. If <code>sig->sigev_notify</code> is <code>SIGEV_SIGNAL</code>, then the signal specified in <code>sig->sigev_signo</code> will be sent to the process. If the <code>SA_SIGINFO</code> flag is set for that signal number, then the signal will be queued to the process and the value specified in <code>sig->sigev_value</code> will be the <code>si_value</code> component of the generated signal (see <code>siginfo(3HEAD)</code>).</p> <p>The <code>list</code> argument is an array of pointers to <code>aiocb</code> structures. The array contains <code>nent</code> elements. The array may contain null elements, which are ignored.</p> <p>The <code>aiocb</code> structure's <code>lio_opcode</code> field of each <code>aiocb</code> structure specifies the operation to be performed. The supported operations are <code>LIO_READ</code>, <code>LIO_WRITE</code>, and <code>LIO_NOP</code>; these symbols are defined in <code><aio.h></code>. The <code>LIO_NOP</code> operation causes the list entry to be ignored. If the <code>aiocb</code> structure's <code>lio_opcode</code> element is equal to <code>LIO_READ</code>, then an I/O operation is submitted as if by a call to <code>aio_read(3RT)</code> with the <code>aiocb</code> equal to the address of the <code>aiocb</code> structure. If the <code>aiocb</code> structure's <code>lio_opcode</code> element is equal to <code>LIO_WRITE</code>, then an I/O operation is submitted as if by a call to <code>aio_write(3RT)</code> with the <code>aiocb</code> equal to the address of the <code>aiocb</code> structure.</p> <p>The <code>aiocb</code> structure's <code>fd</code> member specifies the file descriptor on which the operation is to be performed.</p> <p>The <code>aiocb</code> structure's <code>buf</code> member specifies the address of the buffer to or from which the data is to be transferred.</p> <p>The <code>aiocb</code> structure's <code>nbytes</code> member specifies the number of bytes of data to be transferred.</p>

lio_listio(3RT)

The members of the *aio*cb structure further describe the I/O operation to be performed, in a manner identical to that of the corresponding *aio*cb structure when used by the `aio_read(3RT)` and `aio_write(3RT)` functions.

The *nent* argument specifies how many elements are members of the list, that is, the length of the array.

The behavior of this function is altered according to the definitions of synchronized I/O data integrity completion and synchronized I/O file integrity completion if synchronized I/O is enabled on the file associated with `aio_fildes`. (see `fcntl(3HEAD)` definitions of `O_DSYNC` and `O_SYNC`.)

For regular files, no data transfer will occur past the offset maximum established in the open file description associated with *aio*cb->`aio_fildes`.

RETURN VALUES

If the *mode* argument has the value `LIO_NOWAIT`, and the I/O operations are successfully queued, `lio_listio()` returns 0; otherwise, it returns -1, and sets `errno` to indicate the error.

If the *mode* argument has the value `LIO_WAIT`, and all the indicated I/O has completed successfully, `lio_listio()` returns 0; otherwise, it returns -1, and sets `errno` to indicate the error.

In either case, the return value only indicates the success or failure of the `lio_listio()` call itself, not the status of the individual I/O requests. In some cases, one or more of the I/O requests contained in the list may fail. Failure of an individual request does not prevent completion of any other individual request. To determine the outcome of each I/O request, the application must examine the error status associated with each *aio*cb control block. Each error status so returned is identical to that returned as a result of an `aio_read(3RT)` or `aio_write(3RT)` function.

ERRORS

The `lio_listio()` function will fail if:

EAGAIN	The resources necessary to queue all the I/O requests were not available. The error status for each request is recorded in the <code>aio_error</code> member of the corresponding <i>aio</i> cb structure, and can be retrieved using <code>aio_error(3RT)</code> .
EAGAIN	The number of entries indicated by <i>nent</i> would cause the system-wide limit <code>AIO_MAX</code> to be exceeded.
EINVAL	The <i>mode</i> argument is an improper value, or the value of <i>nent</i> is greater than <code>AIO_LISTIO_MAX</code> .
EINTR	A signal was delivered while waiting for all I/O requests to complete during an <code>LIO_WAIT</code> operation. Note that, since each I/O operation invoked by <code>lio_listio()</code> may possibly provoke a signal when it completes, this error return may be caused by the completion of one (or more) of the very I/O operations being awaited. Outstanding I/O requests are not canceled, and the

application can use `aio_fsync(3RT)` to determine if any request was initiated; `aio_return(3RT)` to determine if any request has completed; or `aio_error(3RT)` to determine if any request was canceled.

EIO One or more of the individual I/O operations failed. The application can use `aio_error(3RT)` to check the error status for each `aio_cb` structure to determine the individual request(s) that failed.

ENOSYS The `lio_listio()` function is not supported by the system.

In addition to the errors returned by the `lio_listio()` function, if the `lio_listio()` function succeeds or fails with errors of `EAGAIN`, `EINTR`, or `EIO`, then some of the I/O specified by the list may have been initiated. If the `lio_listio()` function fails with an error code other than `EAGAIN`, `EINTR`, or `EIO`, no operations from the list have been initiated. The I/O operation indicated by each list element can encounter errors specific to the individual read or write function being performed. In this event, the error status for each `aio_cb` control block contains the associated error code. The error codes that can be set are the same as would be set by a `read(2)` or `write(2)` function, with the following additional error codes possible:

EAGAIN The requested I/O operation was not queued due to resource limitations.

ECANCELED The requested I/O was canceled before the I/O completed due to an explicit `aio_cancel(3RT)` request.

EFBIG The `aio_cb->aio_lio_opcode` is `LIO_WRITE`, the file is a regular file, `aio_cb->aio_nbytes` is greater than 0, and the `aio_cb->aio_offset` is greater than or equal to the offset maximum in the open file description associated with `aio_cb->aio_fildes`.

EINPROGRESS The requested I/O is in progress.

EOVERFLOW The `aio_cb->aio_lio_opcode` is `LIO_READ`, the file is a regular file, `aio_cb->aio_nbytes` is greater than 0, and the `aio_cb->aio_offset` is before the end-of-file and is greater than or equal to the offset maximum in the open file description associated with `aio_cb->aio_fildes`.

USAGE The `lio_listio()` function has a transitional interface for 64-bit file offsets. See `lf64(5)`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

lio_listio(3RT)

SEE ALSO | close(2), exec(2), exit(2), fork(2), lseek(2), read(2), write(2),
aio_cancel(3RT), aio_fsync(3RT), aio_read(3RT), aio_return(3RT),
attributes(5), aio(3HEAD), fcntl(3HEAD), lf64(5), siginfo(3HEAD),
signal(3HEAD)

NOTES | Solaris 2.6 was the first release to support the Asynchronous Input and Output option.
Prior to this release, this function always returned `-1` and set `errno` to `ENOSYS`.

NAME	mq_close – close a message queue				
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <mqqueue.h> int mq_close(mqd_t <i>mqdes</i>);</pre>				
DESCRIPTION	<p>The <code>mq_close()</code> function removes the association between the message queue descriptor, <i>mqdes</i>, and its message queue. The results of using this message queue descriptor after successful return from this <code>mq_close()</code>, and until the return of this message queue descriptor from a subsequent <code>mq_open(3RT)</code>, are undefined.</p> <p>If the process (or thread) has successfully attached a notification request to the message queue via this <i>mqdes</i>, this attachment is removed and the message queue is available for another process to attach for notification.</p>				
RETURN VALUES	Upon successful completion, <code>mq_close()</code> returns 0; otherwise, the function returns -1 and sets <code>errno</code> to indicate the error condition.				
ERRORS	<p>The <code>mq_close()</code> function will fail if:</p> <p>EBADF The <i>mqdes</i> argument is an invalid message queue descriptor.</p> <p>ENOSYS The <code>mq_open()</code> function is not supported by the system.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>mq_notify(3RT)</code> , <code>mq_open(3RT)</code> , <code>mq_unlink(3RT)</code> , <code>attributes(5)</code> , <code>mqqueue(3HEAD)</code>				
NOTES	Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set <code>errno</code> to <code>ENOSYS</code> .				

mq_getattr(3RT)

NAME	mq_getattr – get message queue attributes				
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <mqqueue.h> int mq_getattr(mqd_t mqdes, struct mq_attr *mqstat);</pre>				
DESCRIPTION	<p>The <i>mqdes</i> argument specifies a message queue descriptor. The <code>mq_getattr()</code> function is used to get status information and attributes of the message queue and the open message queue description associated with the message queue descriptor. The results are returned in the <i>mq_attr</i> structure referenced by the <i>mqstat</i> argument.</p> <p>Upon return, the following members will have the values associated with the open message queue description as set when the message queue was opened and as modified by subsequent <code>mq_setattr(3RT)</code> calls:</p> <p><code>mq_flags</code> message queue flags</p> <p>The following attributes of the message queue are returned as set at message queue creation:</p> <p><code>mq_maxmsg</code> maximum number of messages <code>mq_msgsize</code> maximum message size <code>mq_curmsgs</code> number of messages currently on the queue.</p>				
RETURN VALUES	Upon successful completion, the <code>mq_getattr()</code> function returns 0. Otherwise, the function returns -1 and sets <code>errno</code> to indicate the error.				
ERRORS	The <code>mq_getattr()</code> function will fail if:				
	<code>EBADF</code> The <i>mqdes</i> argument is not a valid message queue descriptor.				
	<code>ENOSYS</code> The <code>mq_getattr()</code> function is not supported by the system.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>msgctl(2)</code> , <code>msgget(2)</code> , <code>msgrcv(2)</code> , <code>msgsnd(2)</code> , <code>mq_open(3RT)</code> , <code>mq_send(3RT)</code> , <code>mq_setattr(3RT)</code> , <code>attributes(5)</code> , <code>mqqueue(3HEAD)</code>				
NOTES	Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set <code>errno</code> to <code>ENOSYS</code> .				

NAME	mq_notify – notify process (or thread) that a message is available on a queue
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <mqqueue.h> int mq_notify(mqd_t <i>mqdes</i>, const struct sigevent *<i>notification</i>);</pre>
DESCRIPTION	<p>The <code>mq_notify()</code> function provides an asynchronous mechanism for processes to receive notice that messages are available in a message queue, rather than synchronously blocking (waiting) in <code>mq_receive(3RT)</code>.</p> <p>If <i>notification</i> is not <code>NULL</code>, this function registers the calling process to be notified of message arrival at an empty message queue associated with the message queue descriptor, <i>mqdes</i>. The notification specified by <i>notification</i> will be sent to the process when the message queue transitions from empty to non-empty. At any time, only one process may be registered for notification by a specific message queue. If the calling process or any other process has already registered for notification of message arrival at the specified message queue, subsequent attempts to register for that message queue will fail.</p> <p>The <i>notification</i> argument points to a structure that defines both the signal to be generated and how the calling process will be notified upon I/O completion. If <i>notification</i>-><code>sigev_notify</code> is <code>SIGEV_NONE</code>, then no signal will be posted upon I/O completion, but the error status and the return status for the operation will be set appropriately. If <i>notification</i>-><code>sigev_notify</code> is <code>SIGEV_SIGNAL</code>, then the signal specified in <i>notification</i>-><code>sigev_signo</code> will be sent to the process. If the <code>SA_SIGINFO</code> flag is set for that signal number, then the signal will be queued to the process and the value specified in <i>notification</i>-><code>sigev_value</code> will be the <code>si_value</code> component of the generated signal (see <code>siginfo(3HEAD)</code>).</p> <p>If <i>notification</i> is <code>NULL</code> and the process is currently registered for notification by the specified message queue, the existing registration is removed. The message queue is then available for future registration.</p> <p>When the notification is sent to the registered process, its registration is removed. The message queue is then be available for registration.</p> <p>If a process has registered for notification of message arrival at a message queue and some processes is blocked in <code>mq_receive(3RT)</code> waiting to receive a message when a message arrives at the queue, the arriving message will be received by the appropriate <code>mq_receive(3RT)</code>, and no notification will be sent to the registered process. The resulting behavior is as if the message queue remains empty, and this notification will not be sent until the next arrival of a message at this queue.</p> <p>Any notification registration is removed if the calling process either closes the message queue or exits.</p>
RETURN VALUES	Upon successful completion, <code>mq_notify()</code> returns 0; otherwise, it returns -1 and sets <code>errno</code> to indicate the error.

mq_notify(3RT)

ERRORS The `mq_notify()` function will fail if:

- EBADF The *mqdes* argument is not a valid message queue descriptor.
- EBUSY A process is already registered for notification by the message queue.
- ENOSYS The `mq_notify()` function is not supported by the system.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `mq_close(3RT)`, `mq_open(3RT)`, `mq_receive(3RT)`, `mq_send(3RT)`, `attributes(5)`, `mqqueue(3HEAD)`, `siginfo(3HEAD)`, `signal(3HEAD)`

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned `-1` and set `errno` to `ENOSYS`.

NAME	mq_open – open a message queue								
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <mqqueue.h> mqd_t mq_open(const char *name, /* unsigned long mode, mq_attr attr */ ...);</pre>								
DESCRIPTION	<p>The <code>mq_open()</code> function establishes the connection between a process and a message queue with a message queue descriptor. It creates an open message queue description that refers to the message queue, and a message queue descriptor that refers to that open message queue description. The message queue descriptor is used by other functions to refer to that message queue.</p> <p>The <i>name</i> argument points to a string naming a message queue. The <i>name</i> argument must conform to the construction rules for a path-name. If <i>name</i> is not the name of an existing message queue and its creation is not requested, <code>mq_open()</code> fails and returns an error. The first character of <i>name</i> must be a slash (/) character and the remaining characters of <i>name</i> cannot include any slash characters. For maximum portability, <i>name</i> should include no more than 14 characters, but this limit is not enforced.</p> <p>The <i>oflag</i> argument requests the desired receive and/or send access to the message queue. The requested access permission to receive messages or send messages is granted if the calling process would be granted read or write access, respectively, to a file with the equivalent permissions.</p> <p>The value of <i>oflag</i> is the bitwise inclusive OR of values from the following list. Applications must specify exactly one of the first three values (access modes) below in the value of <i>oflag</i>:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">O_RDONLY</td> <td>Open the message queue for receiving messages. The process can use the returned message queue descriptor with <code>mq_receive(3RT)</code>, but not <code>mq_send(3RT)</code>. A message queue may be open multiple times in the same or different processes for receiving messages.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">O_WRONLY</td> <td>Open the queue for sending messages. The process can use the returned message queue descriptor with <code>mq_send(3RT)</code> but not <code>mq_receive(3RT)</code>. A message queue may be open multiple times in the same or different processes for sending messages.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">O_RDWR</td> <td>Open the queue for both receiving and sending messages. The process can use any of the functions allowed for O_RDONLY and O_WRONLY. A message queue may be open multiple times in the same or different processes for sending messages.</td> </tr> </table> <p>Any combination of the remaining flags may additionally be specified in the value of <i>oflag</i>:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">O_CREAT</td> <td>This option is used to create a message queue, and it requires two additional arguments: <i>mode</i>, which is of type <code>mode_t</code>, and <i>attr</i>,</td> </tr> </table>	O_RDONLY	Open the message queue for receiving messages. The process can use the returned message queue descriptor with <code>mq_receive(3RT)</code> , but not <code>mq_send(3RT)</code> . A message queue may be open multiple times in the same or different processes for receiving messages.	O_WRONLY	Open the queue for sending messages. The process can use the returned message queue descriptor with <code>mq_send(3RT)</code> but not <code>mq_receive(3RT)</code> . A message queue may be open multiple times in the same or different processes for sending messages.	O_RDWR	Open the queue for both receiving and sending messages. The process can use any of the functions allowed for O_RDONLY and O_WRONLY. A message queue may be open multiple times in the same or different processes for sending messages.	O_CREAT	This option is used to create a message queue, and it requires two additional arguments: <i>mode</i> , which is of type <code>mode_t</code> , and <i>attr</i> ,
O_RDONLY	Open the message queue for receiving messages. The process can use the returned message queue descriptor with <code>mq_receive(3RT)</code> , but not <code>mq_send(3RT)</code> . A message queue may be open multiple times in the same or different processes for receiving messages.								
O_WRONLY	Open the queue for sending messages. The process can use the returned message queue descriptor with <code>mq_send(3RT)</code> but not <code>mq_receive(3RT)</code> . A message queue may be open multiple times in the same or different processes for sending messages.								
O_RDWR	Open the queue for both receiving and sending messages. The process can use any of the functions allowed for O_RDONLY and O_WRONLY. A message queue may be open multiple times in the same or different processes for sending messages.								
O_CREAT	This option is used to create a message queue, and it requires two additional arguments: <i>mode</i> , which is of type <code>mode_t</code> , and <i>attr</i> ,								

mq_open(3RT)

which is pointer to a `mq_attr` structure. If the pathname, *name*, has already been used to create a message queue that still exists, then this flag has no effect, except as noted under `O_EXCL` (see below). Otherwise, a message queue is created without any messages in it.

The user ID of the message queue is set to the effective user ID of process, and the group ID of the message queue is set to the effective group ID of the process. The file permission bits are set to the value of *mode*, and modified by clearing all bits set in the file mode creation mask of the process (see `umask(2)`).

If *attr* is non-NULL and the calling process has the appropriate privilege on *name*, the message queue *mq_maxmsg* and *mq_msgsize* attributes are set to the values of the corresponding members in the `mq_attr` structure referred to by *attr*. If *attr* is non-NULL, but the calling process does not have the appropriate privilege on *name*, the `mq_open()` function fails and returns an error without creating the message queue.

`O_EXCL` If both `O_EXCL` and `O_CREAT` are set, `mq_open()` will fail if the message queue *name* exists. The check for the existence of the message queue and the creation of the message queue if it does not exist are atomic with respect to other processes executing `mq_open()` naming the same *name* with both `O_EXCL` and `O_CREAT` set. If `O_EXCL` and `O_CREAT` are not set, the result is undefined.

`O_NONBLOCK` The setting of this flag is associated with the open message queue description and determines whether a `mq_send(3RT)` or `mq_receive(3RT)` waits for resources or messages that are not currently available, or fails with `errno` set to `EAGAIN`. See `mq_send(3RT)` and `mq_receive(3RT)` for details.

RETURN VALUES Upon successful completion, `mq_open()` returns a message queue descriptor; otherwise the function returns `(mqd_t)-1` and sets `errno` to indicate the error condition.

ERRORS The `mq_open()` function will fail if:

<code>EACCESS</code>	The message queue exists and the permissions specified by <i>oflag</i> are denied, or the message queue does not exist and permission to create the message queue is denied.
<code>EEXIST</code>	<code>O_CREAT</code> and <code>O_EXCL</code> are set and the named message queue already exists.
<code>EINTR</code>	The <code>mq_open()</code> operation was interrupted by a signal.

mq_open(3RT)

EINVAL	The <code>mq_open()</code> operation is not supported for the given name, or <code>O_CREAT</code> was specified in <i>oflag</i> , the value of <i>attr</i> is not <code>NULL</code> , and either <code>mq_maxmsg</code> or <code>mq_msgsize</code> was less than or equal to zero.
EMFILE	The number of open message queue descriptors in this process exceeds <code>MQ_OPEN_MAX</code> , or the number of open file descriptors in this process exceeds <code>OPEN_MAX</code> .
ENAMETOOLONG	The length of the <i>name</i> string exceeds <code>PATH_MAX</code> , or a pathname component is longer than <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENFILE	Too many message queues are currently open in the system.
ENOENT	<code>O_CREAT</code> is not set and the named message queue does not exist.
ENOSPC	There is insufficient space for the creation of the new message queue.
ENOSYS	The <code>mq_open()</code> function is not supported by the system.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `exec(2)`, `exit(2)`, `umask(2)`, `mq_close(3RT)`, `mq_receive(3RT)`, `mq_send(3RT)`, `mq_setattr(3RT)`, `mq_unlink(3RT)`, `sysconf(3C)`, `attributes(5)`, `mqueue(3HEAD)`

NOTES Due to the manner in which message queues are implemented, they should not be considered secure and should not be used in security-sensitive applications.

Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned `-1` and set `errno` to `ENOSYS`.

mq_receive(3RT)

NAME	mq_receive – receive a message from a message queue												
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <mqqueue.h> ssize_t mq_receive(mqd_t mqdes, char *msg_ptr, size_t msg_len, unsigned int *msg_prio);</pre>												
DESCRIPTION	<p>The <code>mq_receive()</code> function is used to receive the oldest of the highest priority message(s) from the message queue specified by <code>mqdes</code>. If the size of the buffer in bytes, specified by <code>msg_len</code>, is less than the <code>mq_msgsize</code> member of the message queue, the function fails and returns an error. Otherwise, the selected message is removed from the queue and copied to the buffer pointed to by <code>msg_ptr</code>.</p> <p>If <code>msg_prio</code> is not NULL, the priority of the selected message is stored in the location referenced by <code>msg_prio</code>.</p> <p>If the specified message queue is empty and <code>O_NONBLOCK</code> is not set in the message queue description associated with <code>mqdes</code>, (see <code>mq_open(3RT)</code> and <code>mq_setattr(3RT)</code>), <code>mq_receive()</code> blocks, waiting until a message is enqueued on the message queue, or until <code>mq_receive()</code> is interrupted by a signal. If more than one process (or thread) is waiting to receive a message when a message arrives at an empty queue, then the process of highest priority that has been waiting the longest is selected to receive the message. If the specified message queue is empty and <code>O_NONBLOCK</code> is set in the message queue description associated with <code>mqdes</code>, no message is removed from the queue, and <code>mq_receive()</code> returns an error.</p>												
RETURN VALUES	Upon successful completion, <code>mq_receive()</code> returns the length of the selected message in bytes and the message is removed from the queue. Otherwise, no message is removed from the queue, the function returns a value of <code>-1</code> , and sets <code>errno</code> to indicate the error condition.												
ERRORS	<p>The <code>mq_receive()</code> function will fail if:</p> <table><tr><td>EAGAIN</td><td><code>O_NONBLOCK</code> was set in the message description associated with <code>mqdes</code>, and the specified message queue is empty.</td></tr><tr><td>EBADF</td><td>The <code>mqdes</code> argument is not a valid message queue descriptor open for reading.</td></tr><tr><td>EMSGSIZE</td><td>The specified message buffer size, <code>msg_len</code>, is less than the message size member of the message queue.</td></tr><tr><td>EINTR</td><td>The <code>mq_receive()</code> function operation was interrupted by a signal.</td></tr><tr><td>ENOSYS</td><td>The <code>mq_receive()</code> function is not supported by the system.</td></tr></table> <p>The <code>mq_receive()</code> function may fail if:</p> <table><tr><td>EBADMSG</td><td>A data corruption problem with the message has been detected.</td></tr></table>	EAGAIN	<code>O_NONBLOCK</code> was set in the message description associated with <code>mqdes</code> , and the specified message queue is empty.	EBADF	The <code>mqdes</code> argument is not a valid message queue descriptor open for reading.	EMSGSIZE	The specified message buffer size, <code>msg_len</code> , is less than the message size member of the message queue.	EINTR	The <code>mq_receive()</code> function operation was interrupted by a signal.	ENOSYS	The <code>mq_receive()</code> function is not supported by the system.	EBADMSG	A data corruption problem with the message has been detected.
EAGAIN	<code>O_NONBLOCK</code> was set in the message description associated with <code>mqdes</code> , and the specified message queue is empty.												
EBADF	The <code>mqdes</code> argument is not a valid message queue descriptor open for reading.												
EMSGSIZE	The specified message buffer size, <code>msg_len</code> , is less than the message size member of the message queue.												
EINTR	The <code>mq_receive()</code> function operation was interrupted by a signal.												
ENOSYS	The <code>mq_receive()</code> function is not supported by the system.												
EBADMSG	A data corruption problem with the message has been detected.												

mq_receive(3RT)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `mq_open(3RT)`, `mq_send(3RT)`, `mq_setattr(3RT)`, `attributes(5)`, `mqueue(3HEAD)`

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned `-1` and set `errno` to `ENOSYS`.

mq_send(3RT)

NAME	mq_send – send a message to a message queue
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <mqqueue.h> int mq_send(mqd_t mqdes, const char *msg_ptr, size_t msg_len, unsigned int msg_prio);</pre>
DESCRIPTION	<p>The <code>mq_send()</code> function adds the message pointed to by the argument <code>msg_ptr</code> to the message queue specified by <code>mqdes</code>. The <code>msg_len</code> argument specifies the length of the message in bytes pointed to by <code>msg_ptr</code>. The value of <code>msg_len</code> is less than or equal to the <code>mq_msgsize</code> attribute of the message queue, or <code>mq_send()</code> fails.</p> <p>If the specified message queue is not full, <code>mq_send()</code> behaves as if the message is inserted into the message queue at the position indicated by the <code>msg_prio</code> argument. A message with a larger numeric value of <code>msg_prio</code> is inserted before messages with lower values of <code>msg_prio</code>. A message will be inserted after other messages in the queue, if any, with equal <code>msg_prio</code>. The value of <code>msg_prio</code> must be greater than zero and less than or equal to <code>MQ_PRIO_MAX</code>.</p> <p>If the specified message queue is full and <code>O_NONBLOCK</code> is not set in the message queue description associated with <code>mqdes</code> (see <code>mq_open(3RT)</code> and <code>mq_setattr(3RT)</code>), <code>mq_send()</code> blocks until space becomes available to enqueue the message, or until <code>mq_send()</code> is interrupted by a signal. If more than one thread is waiting to send when space becomes available in the message queue, then the thread of the highest priority which has been waiting the longest is unblocked to send its message. Otherwise, it is unspecified which waiting thread is unblocked. If the specified message queue is full and <code>O_NONBLOCK</code> is set in the message queue description associated with <code>mqdes</code>, the message is not queued and <code>mq_send()</code> returns an error.</p>
RETURN VALUES	Upon successful completion, <code>mq_send()</code> returns 0; otherwise, no message is enqueued, the function returns -1, and <code>errno</code> is set to indicate the error.
ERRORS	The <code>mq_send()</code> function will fail if:
EAGAIN	The <code>O_NONBLOCK</code> flag is set in the message queue description associated with <code>mqdes</code> , and the specified message queue is full.
EBADF	The <code>mqdes</code> argument is not a valid message queue descriptor open for writing.
EINTR	A signal interrupted the call to <code>mq_send()</code>
EINVAL	The value of <code>msg_prio</code> was outside the valid range.
EMSGSIZE	The specified message length, <code>msg_len</code> , exceeds the message size attribute of the message queue.
ENOSYS	The <code>mq_send()</code> function is not supported by the system.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

mq_send(3RT)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO mq_open(3RT), mq_receive(3RT), mq_setattr(3RT), sysconf(3C), attributes(5), mqueue(3HEAD)

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set errno to ENOSYS.

mq_setattr(3RT)

NAME	mq_setattr – set/get message queue attributes				
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <mqeue.h> int mq_setattr(mqd_t mqdes, const struct mq_attr *mqstat, struct mq_attr *omqstat);</pre>				
DESCRIPTION	<p>The <code>mq_setattr()</code> function is used to set attributes associated with the open message queue description referenced by the message queue descriptor specified by <code>mqdes</code>.</p> <p>The message queue attributes corresponding to the following members defined in the <code>mq_attr</code> structure are set to the specified values upon successful completion of <code>mq_setattr()</code>:</p> <p><code>mq_flags</code> The value of this member is either 0 or <code>O_NONBLOCK</code>.</p> <p>The values of <code>mq_maxmsg</code>, <code>mq_msgsize</code>, and <code>mq_curmsgs</code> are ignored by <code>mq_setattr()</code>.</p> <p>If <code>omqstat</code> is non-NULL, <code>mq_setattr()</code> stores, in the location referenced by <code>omqstat</code>, the previous message queue attributes and the current queue status. These values are the same as would be returned by a call to <code>mq_getattr()</code> at that point.</p>				
RETURN VALUES	Upon successful completion, <code>mq_setattr()</code> returns 0 and the attributes of the message queue will have been changed as specified. Otherwise, the message queue attributes are unchanged, and the function returns -1 and sets <code>errno</code> to indicate the error.				
ERRORS	The <code>mq_setattr()</code> function will fail if:				
	<p><code>EBADF</code> The <code>mqdes</code> argument is not a valid message queue descriptor.</p> <p><code>ENOSYS</code> The <code>mq_setattr()</code> function is not supported by the system.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>mq_getattr(3RT)</code> , <code>mq_open(3RT)</code> , <code>mq_receive(3RT)</code> , <code>mq_send(3RT)</code> , <code>attributes(5)</code> , <code>mqeue(3HEAD)</code>				
NOTES	Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set <code>errno</code> to <code>ENOSYS</code> .				

NAME	mq_unlink – remove a message queue								
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <mqueue.h> int mq_unlink(const char *name);</pre>								
DESCRIPTION	<p>The mq_unlink() function removes the message queue named by the pathname <i>name</i>. After a successful call to mq_unlink() with <i>name</i>, a call to mq_open(3RT) with <i>name</i> fails if the flag O_CREAT is not set in <i>flags</i>. If one or more processes have the message queue open when mq_unlink() is called, destruction of the message queue is postponed until all references to the message queue have been closed. Calls to mq_open(3RT) to re-create the message queue may fail until the message queue is actually removed. However, the mq_unlink() call need not block until all references have been closed; it may return immediately.</p>								
RETURN VALUES	<p>Upon successful completion, mq_unlink() returns 0; otherwise, the named message queue is not changed by this function call, the function returns -1 and sets errno to indicate the error.</p>								
ERRORS	<p>The mq_unlink() function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EACCESS</td> <td>Permission is denied to unlink the named message queue.</td> </tr> <tr> <td>ENAMETOOLONG</td> <td>The length of the <i>name</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX while _POSIX_NO_TRUNC is in effect.</td> </tr> <tr> <td>ENOENT</td> <td>The named message queue, <i>name</i>, does not exist.</td> </tr> <tr> <td>ENOSYS</td> <td>mq_unlink() is not supported by the system.</td> </tr> </table>	EACCESS	Permission is denied to unlink the named message queue.	ENAMETOOLONG	The length of the <i>name</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX while _POSIX_NO_TRUNC is in effect.	ENOENT	The named message queue, <i>name</i> , does not exist.	ENOSYS	mq_unlink() is not supported by the system.
EACCESS	Permission is denied to unlink the named message queue.								
ENAMETOOLONG	The length of the <i>name</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX while _POSIX_NO_TRUNC is in effect.								
ENOENT	The named message queue, <i>name</i> , does not exist.								
ENOSYS	mq_unlink() is not supported by the system.								
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
MT-Level	MT-Safe								
SEE ALSO	mq_close(3RT), mq_open(3RT), attributes(5), mqueue(3HEAD)								
NOTES	<p>Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set errno to ENOSYS.</p>								

mutex(3THR)

NAME mutex – concepts relating to mutual exclusion locks

DESCRIPTION Mutual exclusion locks (mutexes) prevent multiple threads from simultaneously executing critical sections of code which access shared data (that is, mutexes are used to serialize the execution of threads). All mutexes must be global. A successful call to acquire a mutex will cause another thread that is also trying to lock the same mutex to block until the owner thread unlocks the mutex.

Mutexes can synchronize threads within the same process or in other processes. Mutexes can be used to synchronize threads between processes if the mutexes are allocated in writable memory and shared among the cooperating processes (see `mmap(2)`), and have been initialized for this task.

The following table lists mutex functions and the actions they perform.

FUNCTION	ACTION
<code>mutex_init</code>	Initialize a mutex.
<code>mutex_destroy</code>	Destroy a mutex.
<code>mutex_lock</code>	Lock a mutex.
<code>mutex_trylock</code>	Attempt to lock a mutex.
<code>mutex_unlock</code>	Unlock a mutex.
<code>pthread_mutex_init</code>	Initialize a mutex.
<code>pthread_mutex_destroy</code>	Destroy a mutex.
<code>pthread_mutex_lock</code>	Lock a mutex.
<code>pthread_mutex_trylock</code>	Attempt to lock a mutex.
<code>pthread_mutex_unlock</code>	Unlock a mutex.

Initialization Mutexes are either intra-process or inter-process, depending upon the argument passed implicitly or explicitly to the initialization of that mutex. A statically allocated mutex does not need to be explicitly initialized; by default, a statically allocated mutex is initialized with all zeros and its scope is set to be within the calling process.

For inter-process synchronization, a mutex needs to be allocated in memory shared between these processes. Since the memory for such a mutex must be allocated dynamically, the mutex needs to be explicitly initialized with the appropriate attribute that indicates inter-process use.

Locking and Unlocking A critical section of code is enclosed by a call to lock the mutex and the call to unlock the mutex to protect it from simultaneous access by multiple threads. Only one thread at a time may possess mutually exclusive access to the critical section of code that is enclosed by the mutex-locking call and the mutex-unlocking call, whether the mutex's

scope is intra-process or inter-process. A thread calling to lock the mutex either gets exclusive access to the code starting from the successful locking until its call to unlock the mutex, or it waits until the mutex is unlocked by the thread that locked it.

Mutexes have ownership, unlike semaphores. Only the thread that locked a mutex, (that is, the owner of the mutex), should unlock it.

If a thread waiting for a mutex receives a signal, upon return from the signal handler, the thread resumes waiting for the mutex as if there was no interrupt.

Caveats Mutexes are almost like data – they can be embedded in data structures, files, dynamic or static memory, and so forth. Hence, they are easy to introduce into a program. However, too many mutexes can degrade performance and scalability of the application. Because too few mutexes can hinder the concurrency of the application, they should be introduced with care. Also, incorrect usage (such as recursive calls, or violation of locking order, and so forth) can lead to deadlocks, or worse, data inconsistencies.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `mmap(2)`, `shmop(2)`, `mutex_destroy(3THR)`, `mutex_init(3THR)`, `mutex_lock(3THR)`, `mutex_trylock(3THR)`, `mutex_unlock(3THR)`, `pthread_mutex_destroy(3THR)`, `pthread_mutex_init(3THR)`, `pthread_mutex_lock(3THR)`, `pthread_mutex_trylock(3THR)`, `pthread_mutex_unlock(3THR)`, `pthread_create(3THR)`, `pthread_mutexattr_init(3THR)`, `attributes(5)`, `standards(5)`

NOTES In the current implementation of threads, `pthread_mutex_lock()`, `pthread_mutex_unlock()`, `mutex_lock()`, `mutex_unlock()`, `pthread_mutex_trylock()`, and `mutex_trylock()` do not validate the mutex type. Therefore, an uninitialized mutex or a mutex with an invalid type does not return `EINVAL`. Interfaces for mutexes with an invalid type have unspecified behavior.

By default, if multiple threads are waiting for a mutex, the order of acquisition is undefined.

`USYNC_THREAD` does not support multiple mappings to the same logical synch object. If you need to `mmap()` a synch object to different locations within the same address space, then the synch object should be initialized as a shared object `USYNC_PROCESS` for Solaris, and `PTHREAD_PROCESS_PRIVATE` for POSIX.

mutex_init(3THR)

NAME	mutex_init, mutex_destroy, mutex_lock, mutex_trylock, mutex_unlock – mutual exclusion locks				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i>[<i>library...</i>] #include <thread.h> #include <synch.h> int mutex_init(mutex_t *mp, int <i>type</i>, void * <i>arg</i>); int mutex_lock(mutex_t *mp); int mutex_trylock(mutex_t *mp); int mutex_unlock(mutex_t *mp); int mutex_destroy(mutex_t *mp);</pre>				
DESCRIPTION	<p>Mutual exclusion locks (mutexes) prevent multiple threads from simultaneously executing critical sections of code which access shared data (that is, mutexes are used to serialize the execution of threads). All mutexes must be global. A successful call for a mutex lock by way of <code>mutex_lock()</code> will cause another thread that is also trying to lock the same mutex to block until the owner thread unlocks it by way of <code>mutex_unlock()</code>. Threads within the same process or within other processes can share mutexes.</p> <p>Mutexes can synchronize threads within the same process or in other processes. Mutexes can be used to synchronize threads between processes if the mutexes are allocated in writable memory and shared among the cooperating processes (see <code>mmap(2)</code>), and have been initialized for this task.</p>				
Initialize	<p>Mutexes are either intra-process or inter-process, depending upon the argument passed implicitly or explicitly to the initialization of that mutex. A statically allocated mutex does not need to be explicitly initialized; by default, a statically allocated mutex is initialized with all zeros and its scope is set to be within the calling process.</p> <p>For inter-process synchronization, a mutex needs to be allocated in memory shared between these processes. Since the memory for such a mutex must be allocated dynamically, the mutex needs to be explicitly initialized using <code>mutex_init()</code>.</p> <p>The <code>mutex_init()</code> function initializes the mutex referenced by <i>mp</i> with the type specified by <i>type</i>. Upon successful initialization the state of the mutex becomes initialized and unlocked. No current type uses <i>arg</i> although a future type may specify additional behavior parameters by way of <i>arg</i>. <i>type</i> may be one of the following:</p> <table><tr><td>USYNC_THREAD</td><td>The mutex can synchronize threads only in this process. <i>arg</i> is ignored.</td></tr><tr><td>USYNC_PROCESS</td><td>The mutex can synchronize threads in this process and other processes. <i>arg</i> is ignored. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see</td></tr></table>	USYNC_THREAD	The mutex can synchronize threads only in this process. <i>arg</i> is ignored.	USYNC_PROCESS	The mutex can synchronize threads in this process and other processes. <i>arg</i> is ignored. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see
USYNC_THREAD	The mutex can synchronize threads only in this process. <i>arg</i> is ignored.				
USYNC_PROCESS	The mutex can synchronize threads in this process and other processes. <i>arg</i> is ignored. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see				

mutex_init(3THR)

mmap(2)). If the object is not allocated in such shared memory, it will not be shared between processes.

USYNC_PROCESS_ROBUST

The mutex can synchronize threads in this process and other processes robustly. At the time of process death, if the lock is held by the process, it is unlocked. The next owner of this mutex will acquire it with an error return of EOWNERDEAD. Note that the application must always check the return code from mutex_lock() for a mutex of this type. The new owner of this mutex should then attempt to make the state protected by the mutex consistent, since this state could have been left inconsistent when the last owner died. If the new owner is able to make the state consistent, it should re-initialize the mutex and then unlock the mutex. If the new owner is not able to make the state consistent, for whatever reason, it should not re-initialize the mutex, but should just unlock the mutex. If the latter event occurs, all waiters will be woken up and all subsequent calls to mutex_lock() will fail in acquiring the mutex with an error code of ENOTRECOVERABLE. mutex can be made consistent by un-initializing the mutex (mutex_destroy()) and re-initializing it (mutex_init()). If the process which got the lock with EOWNERDEAD died, the next owner will get the lock with an error return of EOWNERDEAD. *arg* is ignored. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see shmop(2)) or in memory mapped to a file (see mmap(2)) and memory must be zeroed before initialization. All the processes interested in the robust lock must call mutex_init() at least once to register robust mutex with the system and potentially initialize it. If the object is not allocated in such shared memory, it will not be shared between processes. If mutex_init() is called on a previously initialized mutex mutex_init() will not re-initialize the mutex.

Initializing mutexes can also be accomplished by allocating in zeroed memory (default), in which case, a type of USYNC_THREAD is assumed. The same mutex must not be simultaneously initialized by multiple threads. A mutex lock must not be re-initialized while in use by other threads. If default mutex attributes are used, the macro DEFAULTMUTEX can be used to initialize mutexes that are statically allocated.

Default mutex initialization (intra-process):

```
mutex_t mp;  
mutex_init(&mp, NULL, NULL);
```

mutex_init(3THR)

```
OR mutex_init(&mp, USYNC_THREAD, NULL);
```

```
OR mutex_t mp = DEFAULTMUTEX;
```

```
OR mutex_t mp;  
mp = calloc(1, sizeof (mutex_t));
```

```
OR mutex_t mp;
```

```
mp = malloc(sizeof (mutex_t));
```

```
memset(mp, 0, sizeof (mutex_t));
```

Customized mutex initialization (inter-process):

```
mutex_init(&mp, USYNC_PROCESS, NULL);
```

Customized mutex initialization (inter-process):

```
mutex_init(&mp, USYNC_PROCESS_ROBUST, NULL);
```

Lock and Unlock

A critical section of code is enclosed by a the call to lock the mutex and the call to unlock the mutex to protect it from simultaneous access by multiple threads. Only one thread at a time may possess mutually exclusive access to the critical section of code that is enclosed by the mutex-locking call and the mutex-unlocking call, whether the mutex's scope is intra-process or inter-process. A thread calling to lock the mutex either gets exclusive access to the code starting from the successful locking until its call to unlock the mutex, or it waits until the mutex is unlocked by the thread that locked it.

Mutexes have ownership, unlike semaphores. Although any thread, within the scope of a mutex, can get an unlocked mutex and lock access to the same critical section of code, only the thread that locked a mutex should unlock it.

If a thread waiting for a mutex receives a signal, upon return from the signal handler, the thread resumes waiting for the mutex as if there was no interrupt. A mutex protects code, not data; therefore, strongly bind a mutex with the data by putting both within the same structure, or at least within the same procedure.

A call to `mutex_lock()` locks the mutex object referenced by `mp`. If the mutex is already locked, the calling thread blocks until the mutex is freed; this will return with the mutex object referenced by `mp` in the locked state with the calling thread as its owner. If the current owner of a mutex tries to relock the mutex, it will result in deadlock.

`mutex_trylock()` is the same as `mutex_lock()`, respectively, except that if the mutex object referenced by `mp` is locked (by any thread, including the current thread), the call returns immediately with an error.

`mutex_unlock()` are called by the owner of the mutex object referenced by `mp` to release it. The mutex must be locked and the calling thread must be the one that last

locked the mutex (the owner). If there are threads blocked on the mutex object referenced by *mp* when `mutex_unlock()` is called, the *mp* is freed, and the scheduling policy will determine which thread gets the mutex. If the calling thread is not the owner of the lock, no error status is returned, and the behavior of the program is undefined.

Destroy `mutex_destroy()` destroys the mutex object referenced by *mp*; the mutex object becomes uninitialized. The space used by the destroyed mutex variable is not freed. It needs to be explicitly reclaimed.

RETURN VALUES If successful, these functions return 0. Otherwise, an error number is returned.

ERRORS These functions may fail if:

EFAULT *mp* points to an illegal address.

The `mutex_init()` function will fail if:

EINVAL The value specified by `type` is invalid.

The `mutex_init()` function will fail for `USYNC_PROCESS_ROBUST` type mutex if:

EBUSY The mutex pointed to by *mp* was already initialized. An attempt to re-initialize a mutex previously initialized, but not yet destroyed.

The `mutex_trylock()` function will fail if:

EBUSY The mutex pointed to by *mp* was already locked.

The `mutex_lock()` or `mutex_trylock()` functions will fail for `USYNC_PROCESS_ROBUST` type mutex if:

EOWNERDEAD The last owner of this mutex died while holding the mutex. This mutex is now owned by the caller. The caller must now attempt to make the state protected by the mutex consistent. If it is able to cleanup the state, then it should re-initialize the mutex (see `mutex_init()`) and unlock the mutex. Subsequent calls to `mutex_lock()` will behave normally, as before. If the caller is not able to cleanup the state, the mutex should not be re-initialized, it should be unlocked. Subsequent calls to `mutex_lock()` will fail to acquire the mutex, with the error code, `ENOTRECOVERABLE`. If the owner who got the lock with `EOWNERDEAD` died, the next owner will get the lock with `EOWNERDEAD`.

ELOCKUNMAPPED The last owner of this mutex unmaped the mutex while holding the mutex. This mutex is now owned by the caller. The caller must now attempt to make the state protected by the mutex consistent. If it is able to

mutex_init(3THR)

cleanup the state, then it should re-initialize the mutex unlock the mutex. See mutex_init(3THR). Subsequent calls to mutex_lock() will behave normally, as before. If the caller is not able to cleanup the state, the mutex should not be re-initialized. Subsequent calls to mutex_lock() will fail to acquire the mutex with the error code, ENOTRECOVERABLE.

ENOTRECOVERABLE

The mutex trying to be acquired is protecting the state that has been left irrecoverable when the mutex's last owner could not make the state protected by the mutex consistent. The mutex has not been acquired. This condition can occur when the lock was previously acquired with EOWNERDEAD or ELOCKUNMAPPED and the owner was not able to clean up the state and unlocked the mutex without making the mutex consistent.

Single Gate

The following example uses one global mutex as a gate-keeper to permit each thread exclusive sequential access to the code within the user-defined function "change_global_data." This type of synchronization will protect the state of shared data, but it also prohibits parallelism.

```
/* cc thisfile.c -pthread */
#define _REENTRANT
#include <stdio.h>
#include <thread.h>
#define NUM_THREADS 12
void *change_global_data(void *); /* for thr_create() */
main(int argc, char * argv[]) {
    int i=0;
    for (i=0; i< NUM_THREADS; i++) {
        thr_create(NULL, 0, change_global_data, NULL, 0, NULL);
    }
    while ((thr_join(NULL, NULL, NULL) == 0));
}

void * change_global_data(void *null){
    static mutex_t Global_mutex;
    static int Global_data = 0;
    mutex_lock(&Global_mutex);
    Global_data++;
    sleep(1);
    printf("%d is global data\n",Global_data);
    mutex_unlock(&Global_mutex);
    return NULL;
}
```

Multiple Instruction Single Data

The previous example, the mutex, the code it owns, and the data it protects was enclosed in one function. The next example uses C++ features to accommodate many functions that use just one mutex to protect one data:


```

/* CC thisfile.c -lthread use C++ to compile*/

#define _REENTRANT
#include <stdlib.h>
#include <stdio.h>
#include <thread.h>
#include <errno.h>
#include <iostream.h>
#define NUM_THREADS 16
void *change_global_data(void *); /* for thr_create() */

class Mutected {
private:
    static mutex_t Global_mutex;
    static int Global_data;
public:
    static int add_to_global_data(void);
    static int subtract_from_global_data(void);
};

int Mutected::Global_data = 0;
mutex_t Mutected::Global_mutex;

int Mutected::add_to_global_data() {
    mutex_lock(&Global_mutex);
    Global_data++;
    mutex_unlock(&Global_mutex);
    return Global_data;
}

int Mutected::subtract_from_global_data() {
    mutex_lock(&Global_mutex);
    Global_data--;
    mutex_unlock(&Global_mutex);
    return Global_data;
}

void
main(int argc, char * argv[]) {
    int i=0;
    for (i=0; i< NUM_THREADS; i++) {
        thr_create(NULL, 0, change_global_data, NULL, 0, NULL);
    }
    while ((thr_join(NULL, NULL, NULL) == 0));
}

void * change_global_data(void *) {
    static int switcher = 0;
    if ((switcher++ % 3) == 0) /* one-in-three threads subtracts */
        cout << Mutected::subtract_from_global_data() << endl;
    else
        cout << Mutected::add_to_global_data() << endl;
    return NULL;
}

```

mutex_init(3THR)

Interprocess Locking

A mutex can protect data that is shared among processes. The mutex would need to be initialized as `USYNC_PROCESS`. One process initializes the process-shared mutex and writes it to a file to be mapped into memory by all cooperating processes (see `mmap(2)`). Afterwards, other independent processes can run the same program (whether concurrently or not) and share mutex-protected data.

```
/* cc thisfile.c -pthread */
/* To execute, run the command line "a.out 0 & a.out 1" */

#define _REENTRANT
#include <sys/types.h>
#include <sys/mman.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <stdio.h>
#include <thread.h>
#define INTERPROCESS_FILE "ipc-sharedfile"
#define NUM_ADDTHREADS 12
#define NUM_SUBTRACTTHREADS 10
#define INCREMENT '0'
#define DECREMENT '1'
typedef struct {
    mutex_t      Interprocess_mutex;
    int          Interprocess_data;
} buffer_t;
buffer_t *buffer;

void *add_interprocess_data(), *subtract_interprocess_data();
void create_shared_memory(), test_argv();
int zeroed[sizeof(buffer_t)];
int ipc_fd, i=0;

void
main(int argc, char * argv[]){
    test_argv(argv[1]);

    switch (*argv[1]) {
    case INCREMENT:
        /* Initializes the process-shared mutex */
        /* Should be run prior to running a DECREMENT process */
        create_shared_memory();
        ipc_fd = open(INTERPROCESS_FILE, O_RDWR);
        buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
            PROT_READ|PROT_WRITE, MAP_SHARED, ipc_fd, 0);
        buffer->Interprocess_data = 0;
        mutex_init(&buffer->Interprocess_mutex, USYNC_PROCESS, 0);
        for (i=0; i< NUM_ADDTHREADS; i++)
            thr_create(NULL, 0, add_interprocess_data, argv[1],
                0, NULL);
        break;

    case DECREMENT:
        /* Should be run after the INCREMENT process has run. */
        while(ipc_fd = open(INTERPROCESS_FILE, O_RDWR)) == -1)
```

```

        sleep(1);
        buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
            PROT_READ|PROT_WRITE, MAP_SHARED, ipc_fd, 0);
        for (i=0; i< NUM_SUBTRACTTHREADS; i++)
            thr_create(NULL, 0, subtract_interprocess_data, argv[1],
                0, NULL);
        break;
    } /* end switch */

    while ((thr_join(NULL,NULL,NULL) == 0));
} /* end main */

void *add_interprocess_data(char argv_1[]){
    mutex_lock(&buffer->Interprocess_mutex);
    buffer->Interprocess_data++;
    sleep(2);
    printf("%d is add-interprocess data, and %c is argv\n",
        buffer->Interprocess_data, argv_1[0]);
    mutex_unlock(&buffer->Interprocess_mutex);
    return NULL;
}

void *subtract_interprocess_data(char argv_1[]) {
    mutex_lock(&buffer->Interprocess_mutex);
    buffer->Interprocess_data--;
    sleep(2);
    printf("%d is subtract-interprocess data, and %c is argv\n",
        buffer->Interprocess_data, argv_1[0]);
    mutex_unlock(&buffer->Interprocess_mutex);
    return NULL;
}

void create_shared_memory(){
    int i;
    ipc_fd = creat(INTERPROCESS_FILE, O_CREAT|O_RDWR );
    for (i=0; i<sizeof(buffer_t); i++){
        zeroed[i] = 0;
        write(ipc_fd, &zeroed[i],2);
    }
    close(ipc_fd);
    chmod(INTERPROCESS_FILE, S_IRWXU|S_IRWXG|S_IRWXO);
}

void test_argv(char argv1[]) {
    if (argv1 == NULL) {
        printf("use 0 as arg1 for initial process\n \
            or use 1 as arg1 for the second process\n");
        exit(NULL);
    }
}

```

mutex_init(3THR)

cooperating processes (see `mmap(2)`). Afterwards, other independent processes can run the same program (whether concurrently or not) and share mutex-protected data.

The following example shows how to use a `USYNC_PROCESS_ROBUST` type mutex.

```
/* cc thisfile.c -pthread */
/* To execute, run the command line "a.out & a.out 1" */
#include <sys/types.h>
#include <sys/mman.h>
#include <fcntl.h>
#include <stdio.h>
#include <thread.h>
#define INTERPROCESS_FILE "ipc-sharedfile"
typedef struct {
    mutex_t    Interprocess_mutex;
    int        Interprocess_data;
} buffer_t;
buffer_t *buffer;
int make_data_consistent();
void create_shared_memory();
int zeroed[sizeof(buffer_t)];
int ipc_fd, i=0;
main(int argc, char * argv[]) {
    int rc;
    if (argc > 1) {
        while((ipc_fd = open(INTERPROCESS_FILE, O_RDWR)) == -1)
            sleep(1);
        buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
            PROT_READ|PROT_WRITE, MAP_SHARED, ipc_fd, 0);
        mutex_init(&buffer->Interprocess_mutex,
            USYNC_PROCESS_ROBUST, 0);
    } else {
        create_shared_memory();
        ipc_fd = open(INTERPROCESS_FILE, O_RDWR);
        buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
            PROT_READ|PROT_WRITE, MAP_SHARED, ipc_fd, 0);
        buffer->Interprocess_data = 0;
        mutex_init(&buffer->Interprocess_mutex,
            USYNC_PROCESS_ROBUST, 0);
    }
    for(;;) {
        rc = mutex_lock(&buffer->Interprocess_mutex);
        switch (rc) {
            case EOWNERDEAD:
                /* lock acquired.
                 * last owner died holding the lock, try to make
                 * the state associated with the mutex consistent.
                 * If so, make the robust lock consistent by
                 * re-initializing it.
                 */
                if (make_data_consistent())
                    mutex_init(&buffer->Interprocess_mutex,
                        USYNC_PROCESS_ROBUST, 0);
                mutex_unlock(&buffer->Interprocess_mutex);
            case ENOTRECOVERABLE:
                /* lock not acquired.
                 * last owner got the mutex with EOWNERDEAD
                */

```

```

        * mutex is not consistent (and data?),
        * so return from here
        */
        exit(1);
        break;
    case 0:
        /* no error - data is consistent */
        /* do something with data */
        mutex_unlock(&buffer->Interprocess_mutex);
        break;
    }
} /* end main */
void create_shared_memory() {
    int i;
    ipc_fd = creat(INTERPROCESS_FILE, O_CREAT|O_RDWR);
    for (i=0; i<sizeof(buffer_t); i++) {
        zeroed[i] = 0;
        write(ipc_fd, &zeroed[i], 2);
    }
    close(ipc_fd);
    chmod(INTERPROCESS_FILE, S_IRWXU|S_IRWXG|S_IRWXO);
}

/* return 1 if able to make data consistent, otherwise 0. */
int make_data_consistent () {
    buffer->Interprocess_data = 0;
    return (1);
}

```

Dynamically Allocated Mutexes

The following example allocates and frees memory in which a mutex is embedded.

```

struct record {
    int field1;
    int field2;
    mutex_t m;
} *r;
r = malloc(sizeof(struct record));
mutex_init(&r->m, USYNC_THREAD, NULL);
/*
 * The fields in this record are accessed concurrently
 * by acquiring the embedded lock.
 */

```

The thread execution in this example is as follows:

Thread 1 executes:

```

...
mutex_lock(&r->m);
r->field1++;
mutex_unlock(&r->m);
...

```

Thread 2 executes:

```

...
mutex_lock(&r->m);
localvar = r->field1;
mutex_unlock(&r->m);
...

```

mutex_init(3THR)

Later, when a thread decides to free the memory pointed to by *r*, the thread should call `mutex_destroy()` on the mutexes in this memory.

In the following example, the main thread can do a `thr_join()` on both of the above threads. If there are no other threads using the memory in *r*, the main thread can now safely free *r*:

```
for (i = 0; i < 2; i++)
    thr_join(0, 0, 0);
mutex_destroy(&r->m); /* first destroy mutex */
free(r);             /* then free memory */
```

If the mutex is not destroyed, the program could have memory leaks.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

`mmap(2)`, `shmop(2)`, `mutex(3THR)`, `attributes(5)`, `standards(5)`

NOTES

Currently, the only supported policy is `SCHED_OTHER`. In Solaris, under the `SCHED_OTHER` policy, there is no established order in which threads are unblocked.

In the current implementation of threads, `mutex_lock()`, `mutex_unlock()`, and `mutex_trylock()` do not validate the mutex type. Therefore, an uninitialized mutex or a mutex with an invalid type does not return `EINVAL`. Interfaces for mutexes with an invalid type have unspecified behavior.

Uninitialized mutexes which are allocated locally may contain junk data. Such mutexes need to be initialized using `mutex_init()`.

By default, if multiple threads are waiting for a mutex, the order of acquisition is undefined.

NAME	nanosleep – high resolution sleep						
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <time.h> int nanosleep(const struct timespec *rqtp, struct timespec *rmtp);</pre>						
DESCRIPTION	<p>The nanosleep() function causes the current thread to be suspended from execution until either the time interval specified by the <i>rqtp</i> argument has elapsed or a signal is delivered to the calling thread and its action is to invoke a signal-catching function or to terminate the process. The suspension time may be longer than requested because the argument value is rounded up to an integer multiple of the sleep resolution or because of the scheduling of other activity by the system. But, except for the case of being interrupted by a signal, the suspension time will not be less than the time specified by <i>rqtp</i>, as measured by the system clock, CLOCK_REALTIME.</p> <p>The use of the nanosleep() function has no effect on the action or blockage of any signal.</p>						
RETURN VALUES	<p>If the nanosleep() function returns because the requested time has elapsed, its return value is 0.</p> <p>If the nanosleep() function returns because it has been interrupted by a signal, the function returns a value of -1 and sets <i>errno</i> to indicate the interruption. If the <i>rmtp</i> argument is non-NULL, the <i>timespec</i> structure referenced by it is updated to contain the amount of time remaining in the interval (the requested time minus the time actually slept). If the <i>rmtp</i> argument is NULL, the remaining time is not returned.</p> <p>If nanosleep() fails, it returns -1 and sets <i>errno</i> to indicate the error.</p>						
ERRORS	<p>The nanosleep() function will fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">EINTR</td> <td>The nanosleep() function was interrupted by a signal.</td> </tr> <tr> <td>EINVAL</td> <td>The <i>rqtp</i> argument specified a nanosecond value less than zero or greater than or equal to 1000 million.</td> </tr> <tr> <td>ENOSYS</td> <td>The nanosleep() function is not supported by this implementation.</td> </tr> </table>	EINTR	The nanosleep() function was interrupted by a signal.	EINVAL	The <i>rqtp</i> argument specified a nanosecond value less than zero or greater than or equal to 1000 million.	ENOSYS	The nanosleep() function is not supported by this implementation.
EINTR	The nanosleep() function was interrupted by a signal.						
EINVAL	The <i>rqtp</i> argument specified a nanosecond value less than zero or greater than or equal to 1000 million.						
ENOSYS	The nanosleep() function is not supported by this implementation.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1" style="width: 100%; border-collapse: collapse; margin-left: 20px;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
MT-Level	MT-Safe						
SEE ALSO	sleep(3C), attributes(5), time(3HEAD)						

proc_service(3PROC)

NAME	proc_service – process service interfaces
SYNOPSIS	<pre>#include <proc_service.h> ps_err_e ps_pdmodel(struct ps_prochandle *ph, int *data_model); ps_err_e ps_pglobal_lookup(struct ps_prochandle *ph, const char *object_name, const char *sym_name, psaddr_t *sym_addr); ps_err_e ps_pglobal_sym(struct ps_prochandle *ph, const char *object_name, const char *sym_name, ps_sym_t *sym); ps_err_e ps_pread(struct ps_prochandle *ph, psaddr_t addr, void *buf, size_t size); ps_err_e ps_pwrite(struct ps_prochandle *ph, psaddr_t addr, const void *buf, size_t size); ps_err_e ps_pdread(struct ps_prochandle *ph, psaddr_t addr, void *buf, size_t size); ps_err_e ps_pdwrite(struct ps_prochandle *ph, psaddr_t addr, const void *buf, size_t size); ps_err_e ps_ptread(struct ps_prochandle *ph, psaddr_t addr, void *buf, size_t size); ps_err_e ps_ptwrite(struct ps_prochandle *ph, psaddr_t addr, const void *buf, size_t size); ps_err_e ps_pstop(struct ps_prochandle *ph); ps_err_e ps_pcontinue(struct ps_prochandle *ph); ps_err_e ps_lstop(struct ps_prochandle *ph, lwpid_t lwpid); ps_err_e ps_lcontinue(struct ps_prochandle *ph, lwpid_t lwpid); ps_err_e ps_lgetregs(struct ps_prochandle *ph, lwpid_t lwpid, pgregset_t gregset); ps_err_e ps_lsetregs(struct ps_prochandle *ph, lwpid_t lwpid, const pgregset_t gregset); ps_err_e ps_lgetfpregs(struct ps_prochandle *ph, lwpid_t lwpid, prfpregset_t *fpregset); ps_err_e ps_lsetfpregs(struct ps_prochandle *ph, lwpid_t lwpid, const prfpregset_t *fpregset); ps_err_e ps_pauxv(struct ps_prochandle *ph, const auxv_t **auxp); ps_err_e ps_kill(struct ps_prochandle *ph, int sig); ps_err_e ps_rolldtoaddr(struct ps_prochandle *ph, lwpid_t lwpid, psaddr_t go_addr, psaddr_t stop_addr);</pre>


```

void ps_plog(const char *fmt);

SPARC ps_err_e ps_lgetxregsize(struct ps_prochandle *ph, lwpid_t lwpid,
    int *xregsize);

ps_err_e ps_lgetxregs(struct ps_prochandle *ph, lwpid_t lwpid,
    caddr_t xregset);

ps_err_e ps_lsetxregs(struct ps_prochandle *ph, lwpid_t lwpid,
    caddr_t xregset);

IA ps_err_e ps_lgetLDT(struct ps_prochandle *ph, lwpid_t lwpid,
    struct ssd *ldt);

```

DESCRIPTION Every program that links `libthread_db` or `librtld_db` must provide a set of process control primitives that will allow `libthread_db` and `librtld_db` to access memory and registers in the target process, to start and to stop the target process, and to look up symbols in the target process. See `libthread_db(3THR)`. For information on `librtld_db`, refer to the *Linker and Libraries Guide*

Refer to the individual reference manual pages that describe these routines for a functional specification that clients of `libthread_db` and `librtld_db` can use to implement this required interface. `<proc_service.h>` lists the C declarations of these routines

FUNCTIONS	Name	Description
	<code>ps_pdmodel()</code>	Returns the data model of the target process.
	<code>ps_pglobal_lookup()</code>	Looks up the symbol in the symbol table of the load object in the target process and returns its address.
	<code>ps_pglobal_sym()</code>	Looks up the symbol in the symbol table of the load object in the target process and returns its symbol table entry.
	<code>ps_pread()</code>	Copies <code>size</code> bytes from the target process to the controlling process.
	<code>ps_pwrite()</code>	Copies <code>size</code> bytes from the controlling process to the target process.
	<code>ps_pdread()</code>	Identical to <code>ps_pread()</code> .
	<code>ps_pdwrite()</code>	Identical to <code>ps_pwrite()</code> .
	<code>ps_ptread()</code>	Identical to <code>ps_pread()</code> .
	<code>ps_ptwrite()</code>	Identical to <code>ps_pwrite()</code> .
	<code>ps_pstop()</code>	Stops the target process.
	<code>ps_pcontinue()</code>	Resumes target process.

proc_service(3PROC)

- ps_lstop() Stops a single lightweight process (LWP) within the target process.
- ps_lcontinue() Resumes a single LWP within the target process.
- ps_lgetregs() Gets the general registers of the LWP.
- ps_lsetregs() Sets the general registers of the LWP.
- ps_lgetfpregs() Gets the LWP's floating point register set.
- ps_lsetfpregs() Sets the LWP's floating point register set.
- ps_pauxv() Returns a pointer to a read-only copy of the target process's auxiliary vector.
- ps_kill() Sends signal to target process.
- ps_lrolltoaddr() Rolls the LWP out of a critical section when the process is stopped.
- ps_plog() Logs a message.
- SPARC** ps_lgetxregsize() Returns the size of the architecture-dependent extra state registers.
- ps_lgetxregs() Gets the extra state registers of the LWP.
- ps_lsetxregs() Sets the extra state registers of the LWP.
- IA** ps_lgetLDT() Reads the local descriptor table of the LWP.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO libthread_db(3THR), attributes(5)

Linker and Libraries Guide

NAME	ps_lgetregs, ps_lsetregs, ps_lgetfpregs, ps_lsetfpregs, ps_lgetxregsize, ps_lgetxregs, ps_lsetxregs – routines that access the target process register in libthread_db
SYNOPSIS	<pre>#include <proc_service.h> ps_err_e ps_lgetregs(struct ps_prochandle *ph, lwpid_t lid, pgregset_t gregset); ps_err_e ps_lsetregs(struct ps_prochandle *ph, lwpid_t lid, static pgregset_t gregset); ps_err_e ps_lgetfpregs(struct ps_prochandle *ph, lwpid_t lid, prfpregset_t *fpregs); ps_err_e ps_lsetfpregs(struct ps_prochandle *ph, lwpid_t lid, static prfpregset_t *fpregs); ps_err_e ps_lgetxregsize(struct ps_prochandle *ph, lwpid_t lid, int *xregsize); ps_err_e ps_lgetxregs(struct ps_prochandle *ph, lwpid_t lid, caddr_t xregset); ps_err_e ps_lsetxregs(struct ps_prochandle *ph, lwpid_t lid, caddr_t xregset);</pre>
DESCRIPTION	<p>ps_lgetregs(), ps_lsetregs(), ps_lgetfpregs(), ps_lsetfpregs(), ps_lgetxregsize(), ps_lgetxregs(), ps_lsetxregs() read and write register sets from lightweight processes (LWPs) within the target process identified by <i>ph</i>.</p> <p>ps_lgetregs() gets the general registers of the LWP identified by <i>lid</i>, and ps_lsetregs() sets them. ps_lgetfpregs() gets the LWP's floating point register set, while ps_lsetfpregs() sets it.</p>
SPARC Only	<p>ps_lgetxregsize(), ps_lgetxregs(), and ps_lsetxregs() are SPARC-specific. They do not need to be defined by a controlling process on non-SPARC architecture.</p> <p>ps_lgetxregsize() returns in <i>*xregsize</i> the size of the architecture-dependent extra state registers. ps_lgetxregs() gets the extra state registers, and ps_lsetxregs() sets them.</p>
RETURN VALUES	<p>PS_OK The call returned successfully.</p> <p>PS_NOFPREGS Floating point registers are neither available for this architecture nor for this process.</p> <p>PS_NOXREGS Extra state registers are not available on this architecture.</p> <p>PS_ERR The function did not return successfully.</p>
ATTRIBUTES	See attributes(5) for description of the following attributes:

ps_lgetregs(3PROC)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO libthread(3LIB), libthread_db(3THR), proc_service(3PROC),
libthread_db(3LIB), threads(3THR), attributes(5)

NAME ps_pglobal_lookup, ps_pglobal_sym – look up a symbol in the symbol table of the load object in the target process

SYNOPSIS

```
#include <proc_service.h>

ps_err_e ps_pglobal_lookup(struct ps_prochandle *ph, const char
    *object_name, const char *sym_name, ps_addr_t *sym_addr);

ps_err_e ps_pglobal_sym(struct ps_prochandle *ph, const char
    *object_name, const char *sym_name, ps_sym_t *sym);
```

DESCRIPTION

ps_pglobal_lookup() looks up the symbol *sym_name* in the symbol table of the load object *object_name* in the target process identified by *ph*. It returns the symbol's value as an address in the target process in **sym_addr*.

ps_pglobal_sym() looks up the symbol *sym_name* in the symbol table of the load object *object_name* in the target process identified by *ph*. It returns the symbol table entry in **sym*. The value in the symbol table entry is the symbol's value as an address in the target process.

RETURN VALUES

PS_OK The call completed successfully.

PS_NOSYM The specified symbol was not found.

PS_ERR The function did not return successfully.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO kill(2), libthread_db(3LIB), libthread_db(3THR), proc_service(3PROC), threads(3THR), attributes(5)

ps_pread(3PROC)

NAME	ps_pread, ps_pwrite, ps_pread, ps_pwrite, ps_ptread, ps_ptwrite – interfaces in libthread_db that target process memory access				
SYNOPSIS	<pre>#include <proc_service.h> ps_err_e ps_pread(struct ps_prochandle *ph, psaddr_t addr, void *buf, size_t size); ps_err_e ps_pwrite(struct ps_prochandle *ph, psaddr_t addr, const void *buf, size_t size); ps_err_e ps_pread(struct ps_prochandle *ph, psaddr_t addr, void *buf, size_t size); ps_err_e ps_pwrite(struct ps_prochandle *ph, psaddr_t addr, const void *buf, size_t size); ps_err_e ps_ptread(struct ps_prochandle *ph, psaddr_t addr, void *buf, size_t size); ps_err_e ps_ptwrite(struct ps_prochandle *ph, psaddr_t addr, const void *buf, size_t size);</pre>				
DESCRIPTION	<p>These routines copy data between the target process's address space and the controlling process. ps_pread() copies <i>size</i> bytes from address <i>addr</i> in the target process into <i>buf</i> in the controlling process. ps_pwrite() is like ps_pread() except that the direction of the copy is reversed; data is copied from the controlling process to the target process.</p> <p>ps_pread() and ps_ptread() behave identically to ps_pread(). ps_pwrite() and ps_ptwrite() behave identically to ps_pwrite(). These functions can be implemented as simple aliases for the corresponding primary functions. They are artifacts of history that must be maintained.</p>				
RETURN VALUES	<p>PS_OK The call returned successfully. <i>size</i> bytes were copied.</p> <p>PS_BADADDR Some part of the address range from <i>addr</i> through <i>addr+size-1</i> is not part of the target process's address space.</p> <p>PS_ERR The function did not return successfully.</p>				
ATTRIBUTES	See attributes(5) for description of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT Level	Safe				
SEE ALSO	libthread(3LIB), libthread_db(3THR), proc_service(3PROC), libthread_db(3LIB), threads(3THR), attributes(5)				

NAME	ps_pstop, ps_pcontinue, ps_lstop, ps_lcontinue, ps_lrolltoaddr, ps_kill – process and LWP control in libthread_db						
SYNOPSIS	<pre>#include <proc_service.h> ps_err_e ps_pstop(struct ps_prochandle *ph); ps_err_e ps_pcontinue(struct ps_prochandle *ph); ps_err_e ps_lstop(struct ps_prochandle *ph, lwpid_t lwpid); ps_err_e ps_lcontinue(struct ps_prochandle *ph, lwpid_t lwpid); ps_err_e ps_lrolltoaddr(struct ps_prochandle *ph, lwpid_t lwpid, psaddr_t go_addr, psaddr_t stop_addr); ps_err_e ps_kill(struct ps_prochandle *ph, int signum);</pre>						
DESCRIPTION	<p>The <code>ps_pstop()</code> function stops the target process identified by <code>ph</code>, while the <code>ps_pcontinue()</code> function allows it to resume.</p> <p>The <code>libthread_db()</code> function uses <code>ps_pstop()</code> to freeze the target process while it is under inspection. Within the scope of any single call from outside <code>libthread_db</code> to a <code>libthread_db</code> routine, <code>libthread_db</code> will call <code>ps_pstop()</code>, at most once. If it does, it will call <code>ps_pcontinue()</code> within the scope of the same routine.</p> <p>The controlling process may already have stopped the target process when it calls <code>libthread_db</code>. In that case, it is not obligated to resume the target process when <code>libthread_db</code> calls <code>ps_pcontinue()</code>. In other words, <code>ps_pstop()</code> is mandatory, while <code>ps_pcontinue()</code> is advisory. After <code>ps_pstop()</code>, the target process must be stopped; after <code>ps_pcontinue()</code>, the target process may be running.</p> <p>The <code>ps_lstop()</code> and <code>ps_lcontinue()</code> functions stop and resume a single lightweight process (LWP) within the target process <code>ph</code>.</p> <p>The <code>ps_lrolltoaddr()</code> function is used to roll an LWP forward out of a critical section when the process is stopped. It is also used to run the <code>libthread_db</code> agent thread on behalf of <code>libthread</code>. The <code>ps_lrolltoaddr()</code> function is always called with the target process stopped, that is, there has been a preceding call to <code>ps_pstop()</code>. The specified LWP must be continued at the address <code>go_addr</code>, or at its current address if <code>go_addr</code> is NULL. It should then be stopped when its execution reaches <code>stop_addr</code>. This routine does not return until the LWP has stopped at <code>stop_addr</code>.</p> <p>The <code>ps_kill()</code> function directs the signal <code>signum</code> to the target process for which the handle is <code>ph</code>. It has the same semantics as <code>kill(2)</code>.</p>						
RETURN VALUES	<table border="0"> <tr> <td style="padding-right: 20px;">PS_OK</td> <td>The call completed successfully. In the case of <code>ps_pstop()</code>, the target process is stopped.</td> </tr> <tr> <td>PS_BADLID</td> <td>For <code>ps_lstop()</code>, <code>ps_lcontinue()</code> and <code>ps_lrolltoaddr()</code>; there is no LWP with id <code>lwpid</code> in the target process.</td> </tr> <tr> <td>PS_ERR</td> <td>The function did not return successfully.</td> </tr> </table>	PS_OK	The call completed successfully. In the case of <code>ps_pstop()</code> , the target process is stopped.	PS_BADLID	For <code>ps_lstop()</code> , <code>ps_lcontinue()</code> and <code>ps_lrolltoaddr()</code> ; there is no LWP with id <code>lwpid</code> in the target process.	PS_ERR	The function did not return successfully.
PS_OK	The call completed successfully. In the case of <code>ps_pstop()</code> , the target process is stopped.						
PS_BADLID	For <code>ps_lstop()</code> , <code>ps_lcontinue()</code> and <code>ps_lrolltoaddr()</code> ; there is no LWP with id <code>lwpid</code> in the target process.						
PS_ERR	The function did not return successfully.						

ps_pstop(3PROC)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO `kill(2)`, `libthread(3LIB)`, `libthread_db(3LIB)`, `libthread_db(3THR)`, `proc_service(3PROC)`, `threads(3THR)`, `attributes(5)`

NAME	pthread_atfork – register fork handlers
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <sys/types.h> #include <unistd.h> int pthread_atfork(void (*prepare) (void), void (*parent) (void), void (*child) (void));</pre>
DESCRIPTION	<p>The pthread_atfork() function declares fork handlers to be called prior to and following fork(2), within the thread that called fork(). The order of calls to pthread_atfork() is significant.</p> <p>Before fork() processing begins, the <i>prepare</i> fork handler is called. The <i>prepare</i> handler is not called if its address is NULL.</p> <p>The <i>parent</i> fork handler is called after fork() processing finishes in the parent process, and the <i>child</i> fork handler is called after fork() processing finishes in the child process. If the address of <i>parent</i> or <i>child</i> is NULL, then its handler is not called.</p> <p>The <i>prepare</i> fork handler is called in LIFO (last-in first-out) order, whereas the <i>parent</i> and <i>child</i> fork handlers are called in FIFO (first-in first-out) order. This calling order allows applications to preserve locking order.</p>
RETURN VALUES	Upon successful completion, pthread_atfork() returns 0. Otherwise, an error number is returned.
ERRORS	<p>The pthread_atfork() function will fail if:</p> <p>ENOMEM Insufficient table space exists to record the fork handler addresses.</p>
USAGE	Solaris threads do not offer pthread_atfork() functionality, though a Solaris threads application can call this interface to ensure fork1()-safety, since the two thread APIs are interoperable. If the application is linked with -lpthread, fork() is defined to be the same as fork1(); otherwise fork() and fork1() behave differently. The pthread_atfork() function affects only fork1(). See fork(2).
EXAMPLES	<p>EXAMPLE 1 make a library safe with respect to fork()</p> <p>All multithreaded applications that call fork() in a POSIX threads program and do more than simply call exec(2) in the child of the fork need to ensure that the child is protected from deadlock.</p> <p>Since the "fork-one" model results in duplicating only the thread that called fork(), it is possible that at the time of the call another thread in the parent owns a lock. This thread is not duplicated in the child, so no thread will unlock this lock in the child. Deadlock occurs if the single thread in the child needs this lock.</p> <p>The problem is more serious with locks in libraries. Since a library writer does not know if the application using the library calls fork(), the library must protect itself from such a deadlock scenario. If the application that links with this library calls</p>

pthread_atfork(3THR)

EXAMPLE 1 make a library safe with respect to `fork()` (Continued)

`fork()` and does not call `exec()` in the child, and if it needs a library lock that may be held by some other thread in the parent that is inside the library at the time of the fork, the application deadlocks inside the library.

The following describes how to make a library safe with respect to `fork()` by using `pthread_atfork()`.

1. Identify all locks used by the library (for example $\{L1, \dots, Ln\}$). Identify also the locking order for these locks (for example $\{L1 \dots Ln\}$, as well.)
2. Add a call to `pthread_atfork(f1, f2, f3)` in the library's `.init` section. `f1`, `f2`, `f3` are defined as follows:

```
f1( )
{
    /* ordered in lock order */
    pthread_mutex_lock(L1);
    pthread_mutex_lock( . . . );
    pthread_mutex_lock(Ln);
}

f2( )
{
    pthread_mutex_unlock(L1);
    pthread_mutex_unlock( . . . );
    pthread_mutex_unlock(Ln);
}

f3( )
{
    pthread_mutex_unlock(L1);
    pthread_mutex_unlock( . . . );
    pthread_mutex_unlock(Ln);
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `exec(2)`, `fork(2)`, `atexit(3C)`, `attributes(5)`, `standards(5)`

pthread_attr_getdetachstate(3THR)

NAME pthread_attr_getdetachstate, pthread_attr_setdetachstate – get or set detachstate attribute

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

int pthread_attr_setdetachstate(pthread_attr_t *attr, int detachstate);

int pthread_attr_getdetachstate(const pthread_attr_t *attr, int *detachstate);
```

DESCRIPTION

The *detachstate* attribute controls whether the thread is created in a detached state. If the thread is created detached, then use of the ID of the newly created thread by the pthread_detach() or pthread_join() function is an error.

The pthread_attr_setdetachstate() and pthread_attr_getdetachstate(), respectively, set and get the *detachstate* attribute in the *attr* object.

The *detachstate* can be set to either PTHREAD_CREATE_DETACHED or PTHREAD_CREATE_JOINABLE. A value of PTHREAD_CREATE_DETACHED causes all threads created with *attr* to be in the detached state, whereas using a value of PTHREAD_CREATE_JOINABLE causes all threads created with *attr* to be in the joinable state. The default value of the *detachstate* attribute is PTHREAD_CREATE_JOINABLE.

RETURN VALUES

Upon successful completion, pthread_attr_setdetachstate() and pthread_attr_getdetachstate() return a value of 0. Otherwise, an error number is returned to indicate the error.

The pthread_attr_getdetachstate() function stores the value of the *detachstate* attribute in *detachstate* if successful.

ERRORS

The pthread_attr_setdetachstate() or pthread_attr_getdetachstate() functions may fail if:

EINVAL *attr* or *detachstate* is invalid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

pthread_attr_init(3THR), pthread_attr_setstackaddr(3THR), pthread_attr_setstacksize(3THR), pthread_create(3THR), attributes(5), standards(5)

pthread_attr_getguardsize(3THR)

NAME	pthread_attr_getguardsize, pthread_attr_setguardsize – get or set the thread guardsize attribute
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_attr_getguardsize(const pthread_attr_t *attr, size_t *guardsize); int pthread_attr_setguardsize(pthread_attr_t *attr, size_t guardsize);</pre>
DESCRIPTION	<p>The <i>guardsize</i> attribute controls the size of the guard area for the created thread's stack. The <i>guardsize</i> attribute provides protection against overflow of the stack pointer. If a thread's stack is created with guard protection, the implementation allocates extra memory at the overflow end of the stack as a buffer against stack overflow of the stack pointer. If an application overflows into this buffer an error results (possibly in a SIGSEGV signal being delivered to the thread).</p> <p>The <i>guardsize</i> attribute is provided to the application for two reasons:</p> <ol style="list-style-type: none">1. Overflow protection can potentially result in wasted system resources. An application that creates a large number of threads, and which knows its threads will never overflow their stack, can save system resources by turning off guard areas.2. When threads allocate large data structures on the stack, large guard areas may be needed to detect stack overflow. <p>The <code>pthread_attr_getguardsize()</code> function gets the <i>guardsize</i> attribute in the <i>attr</i> object. This attribute is returned in the <i>guardsize</i> parameter.</p> <p>The <code>pthread_attr_setguardsize()</code> function sets the <i>guardsize</i> attribute in the <i>attr</i> object. The new value of this attribute is obtained from the <i>guardsize</i> parameter. If <i>guardsize</i> is 0, a guard area will not be provided for threads created with <i>attr</i>. If <i>guardsize</i> is greater than 0, a guard area of at least size <i>guardsize</i> bytes is provided for each thread created with <i>attr</i>.</p> <p>A conforming implementation is permitted to round up the value contained in <i>guardsize</i> to a multiple of the configurable system variable <code>PAGESIZE</code>. If an implementation rounds up the value of <i>guardsize</i> to a multiple of <code>PAGESIZE</code>, a call to <code>pthread_attr_getguardsize()</code> specifying <i>attr</i> will store in the <i>guardsize</i> parameter the guard size specified by the previous <code>pthread_attr_setguardsize()</code> function call.</p> <p>The default value of the <i>guardsize</i> attribute is <code>PAGESIZE</code> bytes. The actual value of <code>PAGESIZE</code> is implementation-dependent and may not be the same on all implementations.</p> <p>If the <i>stackaddr</i> attribute has been set (that is, the caller is allocating and managing its own thread stacks), the <i>guardsize</i> attribute is ignored and no protection will be</p>

pthread_attr_getguardsize(3THR)

provided by the implementation. It is the responsibility of the application to manage stack overflow along with stack allocation and management in this case.

RETURN VALUES If successful, the `pthread_attr_getguardsize()` and `pthread_attr_setguardsize()` functions return 0. Otherwise, an error number is returned to indicate the error.

ERRORS The `pthread_attr_getguardsize()` and `pthread_attr_setguardsize()` functions will fail if:

- EINVAL The attribute *attr* is invalid.
- EINVAL The parameter *guardsize* is invalid.
- EINVAL The parameter *guardsize* contains an invalid value.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `sysconf(3C)`, `pthread_attr_init(3THR)`, `attributes(5)`

pthread_attr_getinheritsched(3THR)

NAME	pthread_attr_getinheritsched, pthread_attr_setinheritsched – get or set inheritsched attribute				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_setinheritsched(pthread_attr_t *attr, int inheritsched); int pthread_attr_getinheritsched(const pthread_attr_t *attr, int *inheritsched);</pre>				
DESCRIPTION	<p>The functions pthread_attr_setinheritsched() and pthread_attr_getinheritsched(), respectively, set and get the <i>inheritsched</i> attribute in the <i>attr</i> argument.</p> <p>When the attribute objects are used by pthread_create(), the <i>inheritsched</i> attribute determines how the other scheduling attributes of the created thread are to be set:</p> <table><tr><td>PTHREAD_INHERIT_SCHED</td><td>Specifies that the scheduling policy and associated attributes are to be inherited from the creating thread, and the scheduling attributes in this <i>attr</i> argument are to be ignored.</td></tr><tr><td>PTHREAD_EXPLICIT_SCHED</td><td>Specifies that the scheduling policy and associated attributes are to be set to the corresponding values from this attribute object.</td></tr></table> <p>The symbols PTHREAD_INHERIT_SCHED and PTHREAD_EXPLICIT_SCHED are defined in the header <pthread.h>.</p>	PTHREAD_INHERIT_SCHED	Specifies that the scheduling policy and associated attributes are to be inherited from the creating thread, and the scheduling attributes in this <i>attr</i> argument are to be ignored.	PTHREAD_EXPLICIT_SCHED	Specifies that the scheduling policy and associated attributes are to be set to the corresponding values from this attribute object.
PTHREAD_INHERIT_SCHED	Specifies that the scheduling policy and associated attributes are to be inherited from the creating thread, and the scheduling attributes in this <i>attr</i> argument are to be ignored.				
PTHREAD_EXPLICIT_SCHED	Specifies that the scheduling policy and associated attributes are to be set to the corresponding values from this attribute object.				
RETURN VALUES	If successful, the pthread_attr_setinheritsched() and pthread_attr_getinheritsched() functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	The pthread_attr_setinheritsched() or pthread_attr_getinheritsched() functions may fail if:	EINVAL	<i>attr</i> or <i>inheritsched</i> is invalid.		
EINVAL	<i>attr</i> or <i>inheritsched</i> is invalid.				
USAGE	After these attributes have been set, a thread can be created with the specified attributes using pthread_create(). Using these routines does not affect the current running thread.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				

pthread_attr_getinheritsched(3THR)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_attr_init(3THR), pthread_attr_setscope(3THR),
pthread_attr_setschedpolicy(3THR),
pthread_attr_setschedparam(3THR), pthread_create(3THR),
pthread_setsched_param(3THR), attributes(5), standards(5)

pthread_attr_getschedparam(3THR)

NAME	pthread_attr_getschedparam, pthread_attr_setschedparam – get or set schedparam attribute				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_setschedparam(pthread_attr_t *attr, const struct sched_param *param); int pthread_attr_getschedparam(const pthread_attr_t *attr, struct sched_param *param);</pre>				
DESCRIPTION	The functions pthread_attr_setschedparam() and pthread_attr_getschedparam(), respectively, set and get the scheduling parameter attributes in the attr argument. The contents of the param structure are defined in <sched.h>. For the SCHED_FIFO and SCHED_RR policies, the only required member of param is sched_priority.				
RETURN VALUES	If successful, the pthread_attr_setschedparam() and pthread_attr_getschedparam() functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	The pthread_attr_setschedparam() function may fail if: EINVAL attr is invalid. The pthread_attr_getschedparam() function may fail if: EINVAL attr or param is invalid.				
USAGE	After these attributes have been set, a thread can be created with the specified attributes using pthread_create(). Using these routines does not affect the current running thread.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	pthread_attr_init(3THR), pthread_attr_setscope(3THR), pthread_attr_setinheritsched(3THR), pthread_attr_setschedpolicy(3THR), pthread_create(3THR), pthread_setschedparam(3THR), attributes(5), standards(5)				

pthread_attr_getschedpolicy(3THR)

NAME pthread_attr_getschedpolicy, pthread_attr_setschedpolicy – get or set schedpolicy attribute

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

int pthread_attr_setschedpolicy(pthread_attr_t *attr, int policy);
int pthread_attr_getschedpolicy(const pthread_attr_t *attr, int
    *policy);
```

DESCRIPTION

The functions pthread_attr_setschedpolicy() and pthread_attr_getschedpolicy(), respectively, set and get the *schedpolicy* attribute in the *attr* argument.

The supported values of *policy* include SCHED_FIFO, SCHED_RR and SCHED_OTHER, which are defined by the header <sched.h>. When threads executing with the scheduling policy SCHED_FIFO or SCHED_RR are waiting on a mutex, they acquire the mutex in priority order when the mutex is unlocked.

RETURN VALUES

If successful, the pthread_attr_setschedpolicy() and pthread_attr_getschedpolicy() functions return 0. Otherwise, an error number is returned to indicate the error.

ERRORS

The pthread_attr_setschedpolicy() or pthread_attr_getschedpolicy() function may fail if:

EINVAL *attr* or *policy* is invalid.

USAGE

After these attributes have been set, a thread can be created with the specified attributes using pthread_create(). Using these routines does not affect the current running thread.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

pthread_attr_init(3THR), pthread_attr_setscope(3THR), pthread_attr_setinheritsched(3THR), pthread_attr_setschedparam(3THR), pthread_create(3THR), pthread_setschedparam(3THR), attributes(5), standards(5)

pthread_attr_getscope(3THR)

NAME	pthread_attr_getscope, pthread_attr_setscope – get or set contention scope attribute				
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_attr_setscope(pthread_attr_t *attr, int contention_scope); int pthread_attr_getscope(const pthread_attr_t *attr, int *contention_scope);</pre>				
DESCRIPTION	<p>The pthread_attr_setscope() and pthread_attr_getscope() functions are used to set and get the <i>contention_scope</i> attribute in the <i>attr</i> object.</p> <p>The <i>contention_scope</i> attribute can have the value PTHREAD_SCOPE_SYSTEM, signifying system scheduling contention scope, or PTHREAD_SCOPE_PROCESS, signifying process scheduling contention scope.</p> <p>The symbols PTHREAD_SCOPE_SYSTEM and PTHREAD_SCOPE_PROCESS are defined by the header <pthread.h>.</p>				
RETURN VALUES	If successful, the pthread_attr_setscope() and pthread_attr_getscope() functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	The pthread_attr_setscope(), or pthread_attr_getscope(), function may fail if:				
	EINVAL <i>attr</i> or <i>contention_scope</i> is invalid.				
USAGE	After these attributes have been set, a thread can be created with the specified attributes using pthread_create(). Using these routines does not affect the current running thread.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	pthread_attr_init(3THR), pthread_attr_setinheritsched(3THR), pthread_attr_setschedpolicy(3THR), pthread_attr_setschedparam(3THR), pthread_create(3THR), pthread_setschedparam(3THR), attributes(5), standards(5)				

pthread_attr_getstackaddr(3THR)

NAME pthread_attr_getstackaddr, pthread_attr_setstackaddr – get or set stackaddr attribute

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

int pthread_attr_setstackaddr(pthread_attr_t *attr, void *stackaddr);

int pthread_attr_getstackaddr(const pthread_attr_t *attr, void
    **stackaddr);
```

DESCRIPTION

The functions pthread_attr_setstackaddr() and pthread_attr_getstackaddr(), respectively, set and get the thread creation stackaddr attribute in the attr object. The stackaddr default is NULL. See pthread_create(3THR).

The stackaddr attribute specifies the location of storage to be used for the created thread's stack. The size of the storage is at least PTHREAD_STACK_MIN.

RETURN VALUES

Upon successful completion, pthread_attr_setstackaddr() and pthread_attr_getstackaddr() return a value of 0. Otherwise, an error number is returned to indicate the error.

If successful, the pthread_attr_getstackaddr() function stores the stackaddr attribute value in stackaddr.

ERRORS

The pthread_attr_setstackaddr() function may fail if:

EINVAL attr is invalid.

The pthread_attr_getstackaddr() function may fail if:

EINVAL attr or stackaddr is invalid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_attr_init(3THR), pthread_attr_setdetachstate(3THR), pthread_attr_setstacksize(3THR), pthread_create(3THR), attributes(5), standards(5)

pthread_attr_getstacksize(3THR)

NAME	pthread_attr_getstacksize, pthread_attr_setstacksize – get or set stacksize attribute				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_setstacksize(pthread_attr_t *attr, size_t stacksize) ; int pthread_attr_getstacksize(const pthread_attr_t *attr, size_t *stacksize) ;</pre>				
DESCRIPTION	<p>The functions pthread_attr_setstacksize() and pthread_attr_getstacksize(), respectively, set and get the thread creation stacksize attribute in the attr object.</p> <p>The stacksize attribute defines the minimum stack size (in bytes) allocated for the created threads stack. When the stacksize argument is NULL, the default stack size becomes 1 megabyte for 32-bit processes and 2 megabytes for 64-bit processes.</p>				
RETURN VALUES	Upon successful completion, pthread_attr_setstacksize() and pthread_attr_getstacksize() return a value of 0. Otherwise, an error number is returned to indicate the error. The pthread_attr_getstacksize() function stores the stacksize attribute value in stacksize if successful.				
ERRORS	The pthread_attr_setstacksize() or pthread_attr_getstacksize() function may fail if: EINVAL attr or stacksize is invalid.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	pthread_attr_init(3THR), pthread_attr_setstackaddr(3THR), pthread_attr_setdetachstate(3THR), pthread_create(3THR), attributes(5), standards(5)				

NAME	pthread_attr_init, pthread_attr_destroy – initialize or destroy threads attribute object																											
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_init(pthread_attr_t *attr); int pthread_attr_destroy(pthread_attr_t *attr);</pre>																											
DESCRIPTION	<p>The function <code>pthread_attr_init()</code> initializes a thread attributes object <i>attr</i> with the default value for all of the individual attributes used by a given implementation.</p> <p>The resulting attribute object (possibly modified by setting individual attribute values), when used by <code>pthread_create()</code>, defines the attributes of the thread created. A single attributes object can be used in multiple simultaneous calls to <code>pthread_create()</code>.</p> <p>The <code>pthread_attr_init()</code> function initializes a thread attributes object (<i>attr</i>) with the default value for each attribute as follows:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Attribute</th> <th style="text-align: left;">Default Value</th> <th style="text-align: left;">Meaning of Default</th> </tr> </thead> <tbody> <tr> <td><i>contentionscope</i></td> <td>PTHREAD_SCOPE_PROCESS</td> <td>resource competition within process</td> </tr> <tr> <td><i>detachstate</i></td> <td>PTHREAD_CREATE_JOINABLE</td> <td>joinable by other threads</td> </tr> <tr> <td><i>stackaddr</i></td> <td>NULL</td> <td>stack allocated by system</td> </tr> <tr> <td><i>stacksize</i></td> <td>NULL</td> <td>1 or 2 megabyte</td> </tr> <tr> <td><i>priority</i></td> <td>0</td> <td>priority of the thread</td> </tr> <tr> <td><i>policy</i></td> <td>SCHED_OTHER</td> <td>determined by system</td> </tr> <tr> <td><i>inheritsched</i></td> <td>PTHREAD_EXPLICIT_SCHED</td> <td>scheduling policy and parameters not inherited but explicitly defined by the attribute object</td> </tr> <tr> <td><i>guardsize</i></td> <td>PAGESIZE</td> <td>size of guard area for a thread's created stack</td> </tr> </tbody> </table> <p>The <code>pthread_attr_destroy()</code> function destroys a thread attributes object (<i>attr</i>), which cannot be reused until it is reinitialized. An implementation may cause <code>pthread_attr_destroy()</code> to set <i>attr</i> to an implementation-dependent invalid value. The behavior of using the attribute after it has been destroyed is undefined.</p>	Attribute	Default Value	Meaning of Default	<i>contentionscope</i>	PTHREAD_SCOPE_PROCESS	resource competition within process	<i>detachstate</i>	PTHREAD_CREATE_JOINABLE	joinable by other threads	<i>stackaddr</i>	NULL	stack allocated by system	<i>stacksize</i>	NULL	1 or 2 megabyte	<i>priority</i>	0	priority of the thread	<i>policy</i>	SCHED_OTHER	determined by system	<i>inheritsched</i>	PTHREAD_EXPLICIT_SCHED	scheduling policy and parameters not inherited but explicitly defined by the attribute object	<i>guardsize</i>	PAGESIZE	size of guard area for a thread's created stack
Attribute	Default Value	Meaning of Default																										
<i>contentionscope</i>	PTHREAD_SCOPE_PROCESS	resource competition within process																										
<i>detachstate</i>	PTHREAD_CREATE_JOINABLE	joinable by other threads																										
<i>stackaddr</i>	NULL	stack allocated by system																										
<i>stacksize</i>	NULL	1 or 2 megabyte																										
<i>priority</i>	0	priority of the thread																										
<i>policy</i>	SCHED_OTHER	determined by system																										
<i>inheritsched</i>	PTHREAD_EXPLICIT_SCHED	scheduling policy and parameters not inherited but explicitly defined by the attribute object																										
<i>guardsize</i>	PAGESIZE	size of guard area for a thread's created stack																										
RETURN VALUES	Upon successful completion, <code>pthread_attr_init()</code> and <code>pthread_attr_destroy()</code> return a value of 0. Otherwise, an error number is returned to indicate the error.																											

pthread_attr_init(3THR)

ERRORS

The `pthread_attr_init()` function will fail if:

`ENOMEM` Insufficient memory exists to initialize the thread attributes object.

The `pthread_attr_destroy()` function may fail if:

`EINVAL` *attr* is invalid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

`sysconf(3C)`, `pthread_attr_getdetachstate(3THR)`,
`pthread_attr_getguardsize(3THR)`,
`pthread_attr_getinheritsched(3THR)`,
`pthread_attr_getschedparam(3THR)`,
`pthread_attr_getschedpolicy(3THR)`, `pthread_attr_getscope(3THR)`,
`pthread_attr_getstackaddr(3THR)`, `pthread_attr_getstacksize(3THR)`,
`pthread_attr_setdetachstate(3THR)`, `pthread_attr_setguardsize(3THR)`,
`pthread_attr_setinheritsched(3THR)`,
`pthread_attr_setschedparam(3THR)`,
`pthread_attr_setschedpolicy(3THR)`, `pthread_attr_setscope(3THR)`,
`pthread_attr_setstackaddr(3THR)`, `pthread_attr_setstacksize(3THR)`,
`pthread_create(3THR)`, `attributes(5)`, `standards(5)`

NAME	pthread_cancel – cancel execution of a thread				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_cancel(pthread_t <i>target_thread</i>);</pre>				
DESCRIPTION	<p>The pthread_cancel() function requests that <i>target_thread</i> be canceled.</p> <p>By default, cancellation is deferred until <i>target_thread</i> reaches a cancellation point. See cancellation(3THR).</p> <p>Cancellation cleanup handlers for <i>target_thread</i> are called when the cancellation is acted on. Upon return of the last cancellation cleanup handler, the thread-specific data destructor functions are called for <i>target_thread</i>. <i>target_thread</i> is terminated when the last destructor function returns.</p> <p>The cancellation processing in <i>target_thread</i> runs asynchronously with respect to the calling thread returning from pthread_cancel().</p>				
RETURN VALUES	If successful, the pthread_cancel() function returns 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The pthread_cancel() function may fail if:</p> <p>ESRCH No thread was found with an ID corresponding to that specified by the given thread ID, <i>target_thread</i>.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	cancellation(3THR), condition(3THR), pthread_cleanup_pop(3THR), pthread_cleanup_push(3THR), pthread_cond_wait(3THR), pthread_cond_timedwait(3THR), pthread_exit(3THR), pthread_join(3THR), pthread_setcancelstate(3THR), pthread_setcanceltype(3THR), pthread_testcancel(3THR), setjmp(3C), attributes(5)				
NOTES	See cancellation(3THR) for a discussion of cancellation concepts.				

pthread_cleanup_pop(3THR)

NAME	pthread_cleanup_pop – pop a thread cancellation cleanup handler				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> void pthread_cleanup_pop(int <i>execute</i>);</pre>				
DESCRIPTION	<p>pthread_cleanup_pop() removes the cleanup handler routine at the top of the cancellation cleanup stack of the calling thread and executes it if <i>execute</i> is non-zero.</p> <p>When the thread calls pthread_cleanup_pop() with a non-zero <i>execute</i> argument, the argument at the top of the stack is popped and executed. An argument of 0 pops the handler without executing it.</p> <p>The Solaris system generates a compile time error if pthread_cleanup_push() does not have a matching pthread_cleanup_pop().</p> <p>Be aware that using longjmp() or siglongjmp() to jump into or out of a push/pop pair can lead to trouble, as either the matching push or the matching pop statement might not get executed.</p>				
RETURN VALUES	The pthread_cleanup_pop() function returns no value.				
ERRORS	No errors are defined.				
	The pthread_cleanup_pop() function will not return an error code of EINTR.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	cancellation(3THR), condition(3THR), pthread_cancel(3THR), pthread_cleanup_push(3THR), pthread_exit(3THR), pthread_join(3THR), pthread_setcancelstate(3THR), pthread_setcanceltype(3THR), pthread_testcancel(3THR), setjmp(3C), attributes(5)				
NOTES	See cancellation(3THR) for a discussion of cancellation concepts.				

pthread_cleanup_push(3THR)

NAME pthread_cleanup_push – push a thread cancellation cleanup handler

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

void pthread_cleanup_push(void (*handler, void *), void *arg);
```

DESCRIPTION pthread_cleanup_push() pushes the specified cancellation cleanup handler routine, *handler*, onto the cancellation cleanup stack of the calling thread.

When a thread exits or is canceled and its cancellation cleanup stack is not empty, the cleanup handlers are invoked with the argument *arg* in last in, first out (LIFO) order from the cancellation cleanup stack.

The Solaris system generates a compile time error if pthread_cleanup_push() does not have a matching pthread_cleanup_pop().

Be aware that using longjmp() or siglongjmp() to jump into or out of a push/pop pair can lead to trouble, as either the matching push or the matching pop statement might not get executed.

RETURN VALUES The pthread_cleanup_push() function returns no value.

ERRORS No errors are defined.

The pthread_cleanup_push() function will not return an error code of EINTR.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO cancellation(3THR), condition(3THR), longjmp(3C), pthread_cancel(3THR), pthread_cleanup_pop(3THR), pthread_exit(3THR), pthread_join(3THR), pthread_setcancelstate(3THR), pthread_setcanceltype(3THR), pthread_testcancel(3THR), attributes(5)

NOTES See cancellation(3THR) for a discussion of cancellation concepts.

pthread_condattr_getpshared(3THR)

NAME	pthread_condattr_getpshared, pthread_condattr_setpshared – get or set the process-shared condition variable attributes
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_condattr_getpshared(const pthread_condattr_t *attr, int *pshared); int pthread_condattr_setpshared(pthread_condattr_t *attr, int pshared);</pre>
DESCRIPTION	<p>The <code>pthread_condattr_getpshared()</code> function obtains the value of the <i>process-shared</i> attribute from the attributes object referenced by <i>attr</i>. The <code>pthread_condattr_setpshared()</code> function is used to set the <i>process-shared</i> attribute in an initialized attributes object referenced by <i>attr</i>.</p> <p>The <i>process-shared</i> attribute is set to <code>PTHREAD_PROCESS_SHARED</code> to permit a condition variable to be operated upon by any thread that has access to the memory where the condition variable is allocated, even if the condition variable is allocated in memory that is shared by multiple processes. If the <i>process-shared</i> attribute is <code>PTHREAD_PROCESS_PRIVATE</code>, the condition variable will only be operated upon by threads created within the same process as the thread that initialized the condition variable; if threads of differing processes attempt to operate on such a condition variable, the behavior is undefined. The default value of the attribute is <code>PTHREAD_PROCESS_PRIVATE</code>.</p> <p>Additional attributes, their default values, and the names of the associated functions to get and set those attribute values are implementation-dependent.</p>
RETURN VALUES	<p>If successful, the <code>pthread_condattr_setpshared()</code> function returns 0. Otherwise, an error number is returned to indicate the error.</p> <p>If successful, the <code>pthread_condattr_getpshared()</code> function returns 0 and stores the value of the <i>process-shared</i> attribute of <i>attr</i> into the object referenced by the <i>pshared</i> parameter. Otherwise, an error number is returned to indicate the error.</p>
ERRORS	<p>The <code>pthread_condattr_getpshared()</code> and <code>pthread_condattr_setpshared()</code> functions may fail if:</p> <p><code>EINVAL</code> The value specified by <i>attr</i> is invalid.</p> <p>The <code>pthread_condattr_setpshared()</code> function will fail if:</p> <p><code>EINVAL</code> The new value specified for the attribute is outside the range of legal values for that attribute.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

pthread_condattr_getpshared(3THR)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_condattr_init(3THR), pthread_create(3THR),
pthread_mutex_init(3THR), pthread_cond_init(3THR), attributes(5)

pthread_condattr_init(3THR)

NAME	pthread_condattr_init, pthread_condattr_destroy – initialize or destroy condition variable attributes object				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_condattr_init(pthread_condattr_t *attr); int pthread_condattr_destroy(pthread_condattr_t *attr);</pre>				
DESCRIPTION	<p>The <code>pthread_condattr_init()</code> function initializes a condition variable attributes object <i>attr</i> with the default value for all of the attributes defined by the implementation.</p> <p>At present, the only attribute available is the scope of condition variables. The default scope of the attribute is <code>PTHREAD_PROCESS_PRIVATE</code>.</p> <p>Attempts to initialize previously initialized condition variable attributes object will leave the storage allocated by the previous initialization unallocated.</p> <p>After a condition variable attributes object has been used to initialize one or more condition variables, any function affecting the attributes object (including destruction) does not affect any previously initialized condition variables.</p> <p>The <code>pthread_condattr_destroy()</code> function destroys a condition variable attributes object; the object becomes, in effect, uninitialized. An implementation may cause <code>pthread_condattr_destroy()</code> to set the object referenced by <i>attr</i> to an invalid value. A destroyed condition variable attributes object can be re-initialized using <code>pthread_condattr_init()</code>; the results of otherwise referencing the object after it has been destroyed are undefined.</p> <p>Additional attributes, their default values, and the names of the associated functions to get and set those attribute values are implementation-dependent.</p>				
RETURN VALUES	If successful, the <code>pthread_condattr_init()</code> and <code>pthread_condattr_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The <code>pthread_condattr_init()</code> function will fail if:</p> <table><tr><td>ENOMEM</td><td>Insufficient memory exists to initialize the condition variable attributes object.</td></tr></table> <p>The <code>pthread_condattr_destroy()</code> function may fail if:</p> <table><tr><td>EINVAL</td><td>The value specified by <i>attr</i> is invalid.</td></tr></table>	ENOMEM	Insufficient memory exists to initialize the condition variable attributes object.	EINVAL	The value specified by <i>attr</i> is invalid.
ENOMEM	Insufficient memory exists to initialize the condition variable attributes object.				
EINVAL	The value specified by <i>attr</i> is invalid.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				

pthread_condattr_init(3THR)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_condattr_getpshared(3THR),
pthread_condattr_setpshared(3THR), pthread_cond_init(3THR),
pthread_create(3THR), pthread_mutex_init(3THR), attributes(5)

pthread_cond_init(3THR)

NAME	pthread_cond_init, pthread_cond_destroy – initialize or destroy condition variables				
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_cond_init(pthread_cond_t *cond, const pthread_condattr_t *attr); int pthread_cond_destroy(pthread_cond_t *cond); pthread_cond_t cond= PTHREAD_COND_INITIALIZER;</pre>				
DESCRIPTION	<p>The function <code>pthread_cond_init()</code> initializes the condition variable referenced by <code>cond</code> with attributes referenced by <code>attr</code>. If <code>attr</code> is <code>NULL</code>, the default condition variable attributes are used; the effect is the same as passing the address of a default condition variable attributes object. See <code>pthread_condattr_init(3THR)</code>. Upon successful initialization, the state of the condition variable becomes initialized.</p> <p>Attempting to initialize an already initialized. condition variable results in undefined behavior.</p> <p>The function <code>pthread_cond_destroy()</code> destroys the given condition variable specified by <code>cond</code>; the object becomes, in effect, uninitialized. An implementation may cause <code>pthread_cond_destroy()</code> to set the object referenced by <code>cond</code> to an invalid value. A destroyed condition variable object can be re-initialized using <code>pthread_cond_init()</code>; the results of otherwise referencing the object after it has been destroyed are undefined.</p> <p>It is safe to destroy an initialized condition variable upon which no threads are currently blocked. Attempting to destroy a condition variable upon which other threads are currently blocked results in undefined behavior.</p> <p>In cases where default condition variable attributes are appropriate, the macro <code>PTHREAD_COND_INITIALIZER</code> can be used to initialize condition variables that are statically allocated. The effect is equivalent to dynamic initialization by a call to <code>pthread_cond_init()</code> with parameter <code>attr</code> specified as <code>NULL</code>, except that no error checks are performed.</p>				
RETURN VALUES	If successful, the <code>pthread_cond_init()</code> and <code>pthread_cond_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error. The <code>EBUSY</code> and <code>EINVAL</code> error checks, if implemented, act as if they were performed immediately at the beginning of processing for the function and caused an error return prior to modifying the state of the condition variable specified by <code>cond</code> .				
ERRORS	<p>The <code>pthread_cond_init()</code> function will fail if:</p> <table><tr><td><code>EAGAIN</code></td><td>The system lacked the necessary resources (other than memory) to initialize another condition variable.</td></tr><tr><td><code>ENOMEM</code></td><td>Insufficient memory exists to initialize the condition variable.</td></tr></table> <p>The <code>pthread_cond_init()</code> function may fail if:</p>	<code>EAGAIN</code>	The system lacked the necessary resources (other than memory) to initialize another condition variable.	<code>ENOMEM</code>	Insufficient memory exists to initialize the condition variable.
<code>EAGAIN</code>	The system lacked the necessary resources (other than memory) to initialize another condition variable.				
<code>ENOMEM</code>	Insufficient memory exists to initialize the condition variable.				

pthread_cond_init(3THR)

EBUSY The implementation has detected an attempt to re-initialize the object referenced by *cond*, a previously initialized, but not yet destroyed, condition variable. This condition is not reported.

EINVAL The value specified by *attr* is invalid. This condition is reported.

The pthread_cond_destroy() function may fail if:

EBUSY The implementation has detected an attempt to destroy the object referenced by *cond* while it is referenced (for example, while being used in a pthread_cond_wait() or pthread_cond_timedwait()) by another thread. This condition is not reported.

EINVAL The value specified by *cond* is invalid. This condition is not reported.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO condition(3THR), pthread_cond_signal(3THR), pthread_cond_broadcast(3THR), pthread_cond_wait(3THR), pthread_cond_timedwait(3THR), pthread_condattr_init(3THR), attributes(5), standards(5)

pthread_cond_signal(3THR)

NAME	pthread_cond_signal, pthread_cond_broadcast – signal or broadcast a condition		
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_cond_signal(pthread_cond_t *cond); int pthread_cond_broadcast(pthread_cond_t *cond);</pre>		
DESCRIPTION	<p>These two functions are used to unblock threads blocked on a condition variable.</p> <p>The pthread_cond_signal() call unblocks at least one of the threads that are blocked on the specified condition variable <i>cond</i> (if any threads are blocked on <i>cond</i>).</p> <p>The pthread_cond_broadcast() call unblocks all threads currently blocked on the specified condition variable <i>cond</i>.</p> <p>If more than one thread is blocked on a condition variable, the scheduling policy determines the order in which threads are unblocked. When each thread unblocked as a result of a pthread_cond_signal() or pthread_cond_broadcast() returns from its call to pthread_cond_wait() or pthread_cond_timedwait(), the thread owns the mutex with which it called pthread_cond_wait() or pthread_cond_timedwait(). The thread(s) that are unblocked contend for the mutex according to the scheduling policy (if applicable), and as if each had called pthread_mutex_lock().</p> <p>The pthread_cond_signal() or pthread_cond_broadcast() functions may be called by a thread whether or not it currently owns the mutex that threads calling pthread_cond_wait() or pthread_cond_timedwait() have associated with the condition variable during their waits; however, if predictable scheduling behavior is required, then that mutex is locked by the thread calling pthread_cond_signal() or pthread_cond_broadcast().</p> <p>The pthread_cond_signal() and pthread_cond_broadcast() functions have no effect if there are no threads currently blocked on <i>cond</i>.</p>		
RETURN VALUES	If successful, the pthread_cond_signal() and pthread_cond_broadcast() functions return 0. Otherwise, an error number is returned to indicate the error.		
ERRORS	The pthread_cond_signal() and pthread_cond_broadcast() function may fail if:		
	<table><tr><td>EINVAL</td><td>The value <i>cond</i> does not refer to an initialized condition variable. This condition is not reported.</td></tr></table>	EINVAL	The value <i>cond</i> does not refer to an initialized condition variable. This condition is not reported.
EINVAL	The value <i>cond</i> does not refer to an initialized condition variable. This condition is not reported.		
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:		

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

pthread_cond_signal(3THR)

SEE ALSO condition(3THR), pthread_cond_init(3THR), pthread_cond_wait(3THR),
pthread_cond_timedwait(3THR), attributes(5), standards(5)

pthread_cond_wait(3THR)

NAME	pthread_cond_wait, pthread_cond_timedwait, pthread_cond_reltimedwait_np – wait on a condition
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_cond_wait(pthread_cond_t *cond, pthread_mutex_t *mutex); int pthread_cond_timedwait(pthread_cond_t *cond, pthread_mutex_t *mutex, const struct timespec *abstime); int pthread_cond_reltimedwait_np(pthread_cond_t *cond, pthread_mutex_t *mutex, const struct timespec *reltime);</pre>
DESCRIPTION	<p>The <code>pthread_cond_wait()</code>, <code>pthread_cond_timedwait()</code>, and <code>pthread_cond_reltimedwait_np()</code> functions are used to block on a condition variable. They are called with <i>mutex</i> locked by the calling thread or undefined behaviour will result.</p> <p>These functions atomically release <i>mutex</i> and cause the calling thread to block on the condition variable <i>cond</i>; atomically here means “atomically with respect to access by another thread to the mutex and then the condition variable”. That is, if another thread is able to acquire the mutex after the about-to-block thread has released it, then a subsequent call to <code>pthread_cond_signal()</code> or <code>pthread_cond_broadcast()</code> in that thread behaves as if it were issued after the about-to-block thread has blocked.</p> <p>Upon successful return, the mutex has been locked and is owned by the calling thread.</p> <p>When using condition variables there is always a boolean predicate, an invariant, associated with each condition wait that must be true before the thread should proceed. Spurious wakeups from the <code>pthread_cond_wait()</code>, <code>pthread_cond_timedwait()</code>, or <code>pthread_cond_reltimedwait_np()</code> functions may occur. Since the return from <code>pthread_cond_wait()</code>, <code>pthread_cond_timedwait()</code>, or <code>pthread_cond_reltimedwait_np()</code> does not imply anything about the value of this predicate, the predicate should always be re-evaluated.</p> <p>The order in which blocked threads are awakened by <code>pthread_cond_signal()</code> or <code>pthread_cond_broadcast()</code> is determined by the scheduling policy. See <code> pthreads(3THR)</code>.</p> <p>The effect of using more than one mutex for concurrent <code>pthread_cond_wait()</code>, <code>pthread_cond_timedwait()</code>, or <code>pthread_cond_reltimedwait_np()</code> operations on the same condition variable will result in undefined behavior.</p> <p>A condition wait (whether timed or not) is a cancellation point. When the cancelability enable state of a thread is set to <code>PTHREAD_CANCEL_DEFERRED</code>, a side effect of acting upon a cancellation request while in a condition wait is that the mutex is re-acquired before calling the first cancellation cleanup handler.</p>

pthread_cond_wait(3THR)

A thread that has been unblocked because it has been canceled while blocked in a call to `pthread_cond_wait()` or `pthread_cond_timedwait()` does not consume any condition signal that may be directed concurrently at the condition variable if there are other threads blocked on the condition variable.

The `pthread_cond_timedwait()` function is the same as `pthread_cond_wait()` except that an error is returned if the absolute time specified by *abstime* passes (that is, system time equals or exceeds *abstime*) before the condition *cond* is signaled or broadcasted, or if the absolute time specified by *abstime* has already been passed at the time of the call. When such time-outs occur, `pthread_cond_timedwait()` will nonetheless release and reacquire the mutex referenced by *mutex*. The function `pthread_cond_timedwait()` is also a cancellation point.

The `pthread_cond_reltimedwait_np()` function is a non-standard extension provided by the Solaris version of pthreads as indicated by the “_np” (non-portable) suffix. The `pthread_cond_reltimedwait_np()` function is the same as `pthread_cond_timedwait()` except that the *reltime* argument specifies a non-negative time relative to the current system time rather than an absolute time. An error value is returned if the relative time passes (that is, system time equals or exceeds the starting system time plus the relative time) before the condition *cond* is signaled or broadcasted. When such timeouts occur, `pthread_cond_reltimedwait_np()` releases and reacquires the mutex referenced by *mutex*. The `pthread_cond_reltimedwait_np()` function is also a cancellation point.

If a signal is delivered to a thread waiting for a condition variable, upon return from the signal handler the thread resumes waiting for the condition variable as if it was not interrupted, or it returns 0 due to spurious wakeup.

RETURN VALUES

Except in the case of `ETIMEDOUT`, all these error checks act as if they were performed immediately at the beginning of processing for the function and cause an error return, in effect, prior to modifying the state of the mutex specified by *mutex* or the condition variable specified by *cond*.

Upon successful completion, 0 is returned. Otherwise, an error value is returned to indicate the error.

ERRORS

The `pthread_cond_timedwait()` function will fail if:

`ETIMEDOUT` The absolute time specified by *abstime* to `pthread_cond_timedwait()` has passed.

The `pthread_cond_reltimedwait_np()` function will fail if:

`EINVAL` The value specified by *reltime* is invalid.

`ETIMEDOUT` The relative time specified by *reltime* to `pthread_cond_reltimedwait_np()` has passed.

pthread_cond_wait(3THR)

The `pthread_cond_wait()` and `pthread_cond_timedwait()` functions may fail if:

- EINVAL The value specified by *abstime* is invalid. This condition is reported.
- EINVAL The value specified by *cond* or *mutex* is invalid. This condition is not reported.
- EINVAL Different mutexes were supplied for concurrent `pthread_cond_wait()` or `pthread_cond_timedwait()`, operations on the same condition variable. This condition is not reported.
- EINVAL The mutex was not owned by the current thread at the time of the call. This condition is not reported.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `condition(3THR)`, `pthread_cond_signal(3THR)`, `pthread_cond_broadcast(3THR)`, `attributes(5)`, `standards(5)`

NAME	pthread_create – create a thread
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_create(pthread_t *<i>thread</i>, const pthread_attr_t *<i>attr</i>, void *(*<i>start_routine</i>, void*), void *<i>arg</i>);</pre>
DESCRIPTION	<p>The <code>pthread_create()</code> function is used to create a new thread, with attributes specified by <i>attr</i>, within a process. If <i>attr</i> is NULL, the default attributes are used. (See <code>pthread_attr_init(3THR)</code>). If the attributes specified by <i>attr</i> are modified later, the thread's attributes are not affected. Upon successful completion, <code>pthread_create()</code> stores the ID of the created thread in the location referenced by <i>thread</i>.</p> <p>The thread is created executing <i>start_routine</i> with <i>arg</i> as its sole argument. If the <i>start_routine</i> returns, the effect is as if there was an implicit call to <code>pthread_exit()</code> using the return value of <i>start_routine</i> as the exit status. Note that the thread in which <code>main()</code> was originally invoked differs from this. When it returns from <code>main()</code>, the effect is as if there was an implicit call to <code>exit()</code> using the return value of <code>main()</code> as the exit status.</p> <p>The signal state of the new thread is initialised as follows:</p> <ul style="list-style-type: none"> ■ The signal mask is inherited from the creating thread. ■ The set of signals pending for the new thread is empty. <p>Default thread creation:</p> <pre>pthread_t tid; void *start_func(void *), *arg; pthread_create(&tid, NULL, start_func, arg);</pre> <p>This would have the same effect as:</p> <pre>pthread_attr_t attr; pthread_attr_init(&attr); /* initialize attr with default attributes */ pthread_create(&tid, &attr, start_func, arg);</pre> <p>User-defined thread creation: To create a thread that is scheduled on a system-wide basis, use:</p> <pre>pthread_attr_init(&attr); /* initialize attr with default attributes */ pthread_attr_setscope(&attr, PTHREAD_SCOPE_SYSTEM); /* system-wide contention */ pthread_create(&tid, &attr, start_func, arg);</pre> <p>To customize the attributes for POSIX threads, see <code>pthread_attr_init(3THR)</code>.</p> <p>A new thread created with <code>pthread_create()</code> uses the stack specified by the <i>stackaddr</i> attribute, and the stack continues for the number of bytes specified by the <i>stacksize</i> attribute. By default, the stack size is 1 megabyte for 32-bit processes and 2 megabyte for 64-bit processes (see <code>pthread_attr_setstacksize(3THR)</code>). If the default is used for both the <i>stackaddr</i> and <i>stacksize</i> attributes, <code>pthread_create()</code></p>

pthread_create(3THR)

creates a stack for the new thread with at least 1 megabyte for 32-bit processes and 2 megabyte for 64-bit processes. (For customizing stack sizes, see NOTES).

If `pthread_create()` fails, no new thread is created and the contents of the location referenced by *thread* are undefined.

RETURN VALUES If successful, the `pthread_create()` function returns 0. Otherwise, an error number is returned to indicate the error.

ERRORS The `pthread_create()` function will fail if:

EAGAIN The system lacked the necessary resources to create another thread, or the system-imposed limit on the total number of threads in a process `PTHREAD_THREADS_MAX` would be exceeded.

EINVAL The value specified by *attr* is invalid.

EPERM The caller does not have appropriate permission to set the required scheduling parameters or scheduling policy.

EXAMPLES **EXAMPLE 1** Example of concurrency with multithreading

The following is an example of concurrency with multithreading. Since POSIX threads and Solaris threads are fully compatible even within the same process, this example uses `pthread_create()` if you execute `a.out 0`, or `thr_create()` if you execute `a.out 1`.

Five threads are created that simultaneously perform a time-consuming function, `sleep(10)`. If the execution of this process is timed, the results will show that all five individual calls to `sleep` for ten-seconds completed in about ten seconds, even on a uniprocessor. If a single-threaded process calls `sleep(10)` five times, the execution time will be about 50-seconds.

The command-line to time this process is:

```
/usr/bin/time a.out 0 (for POSIX threading)
```

or

```
/usr/bin/time a.out 1 (for Solaris threading)
```

```
/* cc thisfile.c -lthread -lpthread */
#define _REENTRANT /* basic 3-lines for threads */
#include <pthread.h>
#include <thread.h>

#define NUM_THREADS 5
#define SLEEP_TIME 10

void *sleeping(void *); /* thread routine */
int i;
thread_t tid[NUM_THREADS]; /* array of thread IDs */
```

EXAMPLE 1 Example of concurrency with multithreading (Continued)

```

int
main(int argc, char *argv[])
{
    if (argc == 1) {
        printf("use 0 as arg1 to use pthread_create( )\n");
        printf("or use 1 as arg1 to use thr_create( )\n");
        return (1);
    }

    switch (*argv[1]) {
    case '0': /* POSIX */
        for ( i = 0; i < NUM_THREADS; i++)
            pthread_create(&tid[i], NULL, sleeping,
                (void *)SLEEP_TIME);
        for ( i = 0; i < NUM_THREADS; i++)
            pthread_join(tid[i], NULL);
        break;

    case '1': /* Solaris */
        for ( i = 0; i < NUM_THREADS; i++)
            thr_create(NULL, 0, sleeping, (void *)SLEEP_TIME, 0,
                &tid[i]);
        while (thr_join(0, NULL, NULL) == 0)
            ;
        break;
    } /* switch */
    printf("main( ) reporting that all %d threads have terminated\n", i);
    return (0);
} /* main */

void *
sleeping(void *arg)
{
    int sleep_time = (int)arg;
    printf("thread %d sleeping %d seconds ...\n", thr_self( ), sleep_time);
    sleep(sleep_time);
    printf("\nthread %d awakening\n", thr_self( ));
    return (NULL);
}

```

EXAMPLE 2 If main() had not waited for the completion of the other threads (using pthread_join(3THR) or thr_join(3THR)), it would have continued to process concurrently until it reached the end of its routine and the entire process would have exited prematurely (see exit(2)).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

pthread_create(3THR)

SEE ALSO | `fork(2)`, `sysconf(3C)`, `pthread_attr_init(3THR)`, `pthread_cancel(3THR)`,
`pthread_exit(3THR)`, `pthread_join(3THR)`, `attributes(5)`, `standards(5)`

NOTES | MT application threads execute independently of each other, thus their relative behavior is unpredictable. Therefore, it is possible for the thread executing `main()` to finish before all other user application threads.

`pthread_join(3THR)`, on the other hand, must specify the terminating thread (IDs) for which it will wait.

A user-specified stack size must be greater than the value `PTHREAD_STACK_MIN`. A minimum stack size may not accommodate the stack frame for the user thread function `start_func`. If a stack size is specified, it must accommodate `start_func` requirements and the functions that it may call in turn, in addition to the minimum requirement.

It is usually very difficult to determine the runtime stack requirements for a thread. `PTHREAD_STACK_MIN` specifies how much stack storage is required to execute a `NULL start_func`. The total runtime requirements for stack storage are dependent on the storage required to do runtime linking, the amount of storage required by library runtimes (as `printf()`) that your thread calls. Since these storage parameters are not known before the program runs, it is best to use default stacks. If you know your runtime requirements or decide to use stacks that are larger than the default, then it makes sense to specify your own stacks.

NAME	pthread_detach – detach a thread				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_detach(pthread_t <i>thread</i>);</pre>				
DESCRIPTION	<p>The <code>pthread_detach()</code> function is used to indicate to the implementation that storage for the thread <i>thread</i> can be reclaimed when that thread terminates. In other words, <code>pthread_detach()</code> dynamically resets the <i>detachstate</i> attribute of the thread to <code>PTHREAD_CREATE_DETACHED</code>. After a successful call to this function, it would not be necessary to reclaim the thread using <code>pthread_join()</code>. See <code>pthread_join(3THR)</code>. If <i>thread</i> has not terminated, <code>pthread_detach()</code> will not cause it to terminate. The effect of multiple <code>pthread_detach()</code> calls on the same target thread is unspecified.</p>				
RETURN VALUES	If successful, <code>pthread_detach()</code> returns 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The <code>pthread_detach()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The implementation has detected that the value specified by <i>thread</i> does not refer to a joinable thread.</td> </tr> <tr> <td>ESRCH</td> <td>No thread could be found corresponding to that specified by the given thread ID.</td> </tr> </table>	EINVAL	The implementation has detected that the value specified by <i>thread</i> does not refer to a joinable thread.	ESRCH	No thread could be found corresponding to that specified by the given thread ID.
EINVAL	The implementation has detected that the value specified by <i>thread</i> does not refer to a joinable thread.				
ESRCH	No thread could be found corresponding to that specified by the given thread ID.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>pthread_create(3THR)</code> , <code>pthread_join(3THR)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

pthread_equal(3THR)

NAME pthread_equal – compare thread IDs

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

int pthread_equal(pthread_t t1, pthread_t t2);
```

DESCRIPTION The pthread_equal() function compares the thread IDs *t1* and *t2*.

RETURN VALUES The pthread_equal() function returns a non-zero value if *t1* and *t2* are equal. Otherwise, 0 is returned.

If *t1* or *t2* is an invalid thread ID, the behavior is undefined.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_create(3THR), pthread_self(3THR), attributes(5)

NOTES Solaris thread IDs do not require an equivalent function because the thread_t structure is an unsigned int.

NAME	pthread_exit – terminate calling thread				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> void pthread_exit(void *<i>value_ptr</i>);</pre>				
DESCRIPTION	<p>The pthread_exit() function terminates the calling thread, in a similar way that exit(3C) terminates the calling process. If the thread is not detached, the exit status specified by <i>value_ptr</i> is made available to any successful join with the terminating thread. See pthread_join(3THR). Any cancellation cleanup handlers that have been pushed and not yet popped are popped in the reverse order that they were pushed and then executed. After all cancellation cleanup handlers have been executed, if the thread has any thread-specific data, appropriate destructor functions will be called in an unspecified order. Thread termination does not release any application visible process resources, including, but not limited to, mutexes and file descriptors, nor does it perform any process level cleanup actions, including, but not limited to, calling any atexit() routines that may exist.</p> <p>An implicit call to pthread_exit() is made when a thread other than the thread in which main() was first invoked returns from the start routine that was used to create it. The function's return value serves as the thread's exit status.</p> <p>The behavior of pthread_exit() is undefined if called from a cancellation cleanup handler or destructor function that was invoked as a result of either an implicit or explicit call to pthread_exit().</p> <p>After a thread has terminated, the result of access to local (auto) variables of the thread is undefined. Thus, references to local variables of the exiting thread should not be used for the pthread_exit() <i>value_ptr</i> parameter value.</p> <p>The process exits with an exit status of 0 after the last thread has been terminated. The behavior is as if the implementation called exit() with a 0 argument at thread termination time.</p>				
RETURN VALUES	The pthread_exit() function cannot return to its caller.				
ERRORS	No errors are defined.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	exit(3C), pthread_cancel(3THR), pthread_create(3THR), pthread_join(3THR), pthread_key_create(3THR), attributes(5), standards(5)				

pthread_getconcurrency(3THR)

NAME	pthread_getconcurrency, pthread_setconcurrency – get or set level of concurrency		
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_getconcurrency(void); int pthread_setconcurrency(int <i>new_level</i>);</pre>		
DESCRIPTION	<p>Unbound threads in a process may or may not be required to be simultaneously active. By default, the threads implementation ensures that a sufficient number of threads are active so that the process can continue to make progress. While this conserves system resources, it may not produce the most effective level of concurrency.</p> <p>The pthread_setconcurrency() function allows an application to inform the threads implementation of its desired concurrency level, <i>new_level</i>. The actual level of concurrency provided by the implementation as a result of this function call is unspecified.</p> <p>If <i>new_level</i> is 0, it causes the implementation to maintain the concurrency level at its discretion as if pthread_setconcurrency() was never called.</p> <p>The pthread_getconcurrency() function returns the value set by a previous call to the pthread_setconcurrency() function. If the pthread_setconcurrency() function was not previously called, this function returns 0 to indicate that the implementation is maintaining the concurrency level.</p> <p>When an application calls pthread_setconcurrency() it is informing the implementation of its desired concurrency level. The implementation uses this as a hint, not a requirement.</p> <p>If an implementation does not support multiplexing of user threads on top of several kernel scheduled entities, the pthread_setconcurrency() and pthread_getconcurrency() functions will be provided for source code compatibility but they will have no effect when called. To maintain the function semantics, the <i>new_level</i> parameter will be saved when pthread_setconcurrency() is called so that a subsequent call to pthread_getconcurrency() returns the same value.</p>		
RETURN VALUES	<p>If successful, the pthread_setconcurrency() function returns 0. Otherwise, an error number is returned to indicate the error.</p> <p>The pthread_getconcurrency() function always returns the concurrency level set by a previous call to pthread_setconcurrency(). If the pthread_setconcurrency() function has never been called, pthread_getconcurrency() returns 0.</p>		
ERRORS	<p>The pthread_setconcurrency() function will fail if:</p> <table><tr><td>EINVAL</td><td>The value specified by <i>new_level</i> is negative.</td></tr></table>	EINVAL	The value specified by <i>new_level</i> is negative.
EINVAL	The value specified by <i>new_level</i> is negative.		

pthread_getconcurrency(3THR)

EAGAIN The value specific by *new_level* would cause a system resource to be exceeded.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_create(3THR), pthread_attr_init(3THR), attributes(5)

pthread_getschedparam(3THR)

NAME	pthread_getschedparam, pthread_setschedparam – access dynamic thread scheduling parameters
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_getschedparam(pthread_t <i>thread</i>, int *<i>policy</i>, struct sched_param *<i>param</i>); int pthread_setschedparam(pthread_t <i>thread</i>, int <i>policy</i>, const struct sched_param *<i>param</i>);</pre>
DESCRIPTION	<p>The pthread_getschedparam() and pthread_setschedparam() functions allow the scheduling policy and scheduling parameters of individual threads within a multithreaded process to be retrieved and set. Supported policies are SCHED_FIFO, SCHED_RR, and SCHED_OTHER. See pthreads(3THR). For SCHED_FIFO, SCHED_RR, and SCHED_OTHER, the affected scheduling parameter is the <i>sched_priority</i> member of the sched_param structure.</p> <p>The pthread_getschedparam() function retrieves the scheduling policy and scheduling parameters for the thread whose thread ID is given by <i>thread</i> and stores those values in <i>policy</i> and <i>param</i>, respectively. The priority value returned from pthread_getschedparam() is the value specified by the most recent pthread_setschedparam() or pthread_create() call affecting the target thread, and reflects any temporary adjustments to its priority as a result of any priority inheritance or ceiling functions. The pthread_setschedparam() function sets the scheduling policy and associated scheduling parameters for the thread whose thread ID is given by <i>thread</i> to the policy and associated parameters provided in <i>policy</i> and <i>param</i>, respectively.</p> <p>If the pthread_setschedparam() function fails, no scheduling parameters will be changed for the target thread.</p>
RETURN VALUES	If successful, the pthread_getschedparam() and pthread_setschedparam() functions return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	The pthread_getschedparam() function may fail if: ESRCH The value specified by <i>thread</i> does not refer to a existing thread. The pthread_setschedparam() function may fail if: EINVAL The value specified by <i>policy</i> or one of the scheduling parameters associated with the scheduling policy <i>policy</i> is invalid. EPERM The caller does not have the appropriate permission to set either the scheduling parameters or the scheduling policy of the specified thread. ESRCH The value specified by <i>thread</i> does not refer to a existing thread.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

pthread_getschedparam(3THR)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_attr_init(3THR), pthreads(3THR), sched_setparam(3RT), sched_getparam(3RT), sched_setscheduler(3RT), sched_getscheduler(3RT) attributes(5), standards(5)

pthread_getspecific(3THR)

NAME	pthread_getspecific, pthread_setspecific – manage thread-specific data				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_setspecific(pthread_key_t <i>key</i>, const void *<i>value</i>); void *pthread_getspecific(pthread_key_t <i>key</i>);</pre>				
DESCRIPTION	<p>The <code>pthread_setspecific()</code> function associates a thread-specific <i>value</i> with a <i>key</i> obtained by way of a previous call to <code>pthread_key_create()</code>. Different threads may bind different values to the same key. These values are typically pointers to blocks of dynamically allocated memory that have been reserved for use by the calling thread.</p> <p>The <code>pthread_getspecific()</code> function returns the value currently bound to the specified <i>key</i> on behalf of the calling thread.</p> <p>The effect of calling <code>pthread_setspecific()</code> or <code>pthread_getspecific()</code> with a <i>key</i> value not obtained from <code>pthread_key_create()</code> or after <i>key</i> has been deleted with <code>pthread_key_delete()</code> is undefined.</p> <p>Both <code>pthread_setspecific()</code> and <code>pthread_getspecific()</code> may be called from a thread-specific data destructor function. However, calling <code>pthread_setspecific()</code> from a destructor may result in lost storage or infinite loops.</p>				
RETURN VALUES	<p>The <code>pthread_getspecific()</code> function returns the thread-specific data value associated with the given <i>key</i>. If no thread-specific data value is associated with <i>key</i>, then the value <code>NULL</code> is returned.</p> <p>Upon successful completion, the <code>pthread_setspecific()</code> function returns 0. Otherwise, an error number is returned to indicate the error.</p>				
ERRORS	<p>The <code>pthread_setspecific()</code> function will fail if:</p> <table><tr><td>ENOMEM</td><td>Insufficient memory exists to associate the value with the key.</td></tr></table> <p>The <code>pthread_setspecific()</code> function may fail if:</p> <table><tr><td>EINVAL</td><td>The key value is invalid.</td></tr></table> <p>The <code>pthread_getspecific()</code> function does not return errors.</p>	ENOMEM	Insufficient memory exists to associate the value with the key.	EINVAL	The key value is invalid.
ENOMEM	Insufficient memory exists to associate the value with the key.				
EINVAL	The key value is invalid.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>pthread_key_create(3THR)</code> <code>attributes(5)</code> , <code>standards(5)</code>				

NAME	pthread_join – wait for thread termination						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_join(pthread_t <i>thread</i>, void **<i>status</i>);</pre>						
DESCRIPTION	<p>The pthread_join() function suspends processing of the calling thread until the target <i>thread</i> completes. <i>thread</i> must be a member of the current process and it cannot be a detached thread. See pthread_create(3THR).</p> <p>If two or more threads wait for the same thread to complete, all will suspend processing until the thread has terminated, and then one thread will return successfully and the others will return with an error of ESRCH. The pthread_join() function will not block processing of the calling thread if the target <i>thread</i> has already terminated.</p> <p>If a pthread_join() call returns successfully with a non-null <i>status</i> argument, the value passed to pthread_exit(3THR) by the terminating thread will be placed in the location referenced by <i>status</i>.</p> <p>If the pthread_join() calling thread is cancelled, then the target <i>thread</i> will remain joinable by pthread_join(). However, the calling thread may set up a cancellation cleanup handler on <i>thread</i> prior to the join call, which may detach the target <i>thread</i> by calling pthread_detach(3THR). See pthread_detach(3THR) and pthread_cancel(3THR).</p>						
RETURN VALUES	If successful, pthread_join() returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<table border="0"> <tr> <td style="vertical-align: top;">EDEADLK</td> <td>A joining deadlock would occur, such as when a thread attempts to wait for itself.</td> </tr> <tr> <td style="vertical-align: top;">EINVAL</td> <td>The thread corresponding to the given thread ID is a detached thread.</td> </tr> <tr> <td style="vertical-align: top;">ESRCH</td> <td>No thread could be found corresponding to the given thread ID.</td> </tr> </table>	EDEADLK	A joining deadlock would occur, such as when a thread attempts to wait for itself.	EINVAL	The thread corresponding to the given thread ID is a detached thread.	ESRCH	No thread could be found corresponding to the given thread ID.
EDEADLK	A joining deadlock would occur, such as when a thread attempts to wait for itself.						
EINVAL	The thread corresponding to the given thread ID is a detached thread.						
ESRCH	No thread could be found corresponding to the given thread ID.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
MT-Level	MT-Safe						
SEE ALSO	wait(2), pthread_cancel(3THR), pthread_create(3THR), pthread_detach(3THR), pthread_exit(3THR), attributes(5), standards(5)						
NOTES	The pthread_join(3THR) function must specify the <i>thread</i> ID for whose termination it will wait.						

pthread_join(3THR)

Calling `pthread_join()` also "detaches" the thread; that is, `pthread_join()` includes the effect of the `pthread_detach()` function. If a thread were to be cancelled when blocked in `pthread_join()`, an explicit detach would have to be performed in the cancellation cleanup handler. The `pthread_detach()` function exists primarily for this purpose.

NAME	pthread_key_create – create thread-specific data key
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_key_create(pthread_key_t *<i>key</i>, void (*<i>destructor</i>, void*));</pre>
DESCRIPTION	<p>This function creates a thread-specific data key visible to all threads in the process. Key values provided by <code>pthread_key_create()</code> are opaque objects used to locate thread-specific data. Although the same key value may be used by different threads, the values bound to the key by <code>pthread_setspecific()</code> are maintained on a per-thread basis and persist for the life of the calling thread.</p> <p>Upon key creation, the value <code>NULL</code> is associated with the new key in all active threads. Upon thread creation, the value <code>NULL</code> is associated with all defined keys in the new thread.</p> <p>An optional destructor function may be associated with each key value. At thread exit, if a key value has a non-<code>NULL</code> destructor pointer, and the thread has a non-<code>NULL</code> value associated with that key, the function pointed to is called with the current associated value as its sole argument. Destructors can be called in any order.</p> <p>If, after all the destructors have been called for all keys with non-<code>NULL</code> values, there are still some keys with non-<code>NULL</code> values, the process will be repeated. If, after at least <code>PTHREAD_DESTRUCTOR_ITERATIONS</code> iterations of destructor calls for outstanding non-<code>NULL</code> values, there are still some keys with non-<code>NULL</code> values, the process is continued, even though this might result in an infinite loop.</p>
RETURN VALUES	If successful, the <code>pthread_key_create()</code> function stores the newly created key value at <i>key</i> and returns 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The <code>pthread_key_create()</code> function will fail if:</p> <p>EAGAIN The system lacked the necessary resources to create another thread-specific data key, or the system-imposed limit on the total number of keys per process <code>PTHREAD_KEYS_MAX</code> has been exceeded.</p> <p>ENOMEM Insufficient memory exists to create the key.</p> <p>The <code>pthread_key_create()</code> function will not return an error code of <code>EINTR</code>.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

pthread_key_create(3THR)

SEE ALSO | pthread_getspecific(3THR), pthread_setspecific(3THR),
pthread_key_delete(3THR), attributes(5), standards(5)

NAME	pthread_key_delete – delete thread-specific data key				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_key_delete(pthread_key_t <i>key</i>);</pre>				
DESCRIPTION	<p>The <code>pthread_key_delete()</code> function deletes a thread-specific data key previously returned by <code>pthread_key_create()</code>. The thread-specific data values associated with <i>key</i> need not be NULL at the time <code>pthread_key_delete()</code> is called. It is the responsibility of the application to free any application storage or perform any cleanup actions for data structures related to the deleted key or associated thread-specific data in any threads; this cleanup can be done either before or after <code>pthread_key_delete()</code> is called. Any attempt to use <i>key</i> following the call to <code>pthread_key_delete()</code> results in undefined behaviour.</p> <p>The <code>pthread_key_delete()</code> function is callable from within destructor functions. No destructor functions will be invoked by <code>pthread_key_delete()</code>. Any destructor function that may have been associated with <i>key</i> will no longer be called upon thread exit.</p>				
RETURN VALUES	If successful, the <code>pthread_key_delete()</code> function returns 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The <code>pthread_key_delete()</code> function may fail if:</p> <p><code>EINVAL</code> The <i>key</i> value is invalid.</p> <p>The <code>pthread_key_delete()</code> function will not return an error code of <code>EINTR</code>.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>pthread_key_create(3THR)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

pthread_kill(3THR)

NAME pthread_kill – send a signal to a thread

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <signal.h>
#include <pthread.h>

int pthread_kill(pthread_t thread, int sig);
```

DESCRIPTION The `pthread_kill()` function is used to request that a signal be delivered to the specified thread.

As in `kill()`, if *sig* is 0, error checking is performed but no signal is actually sent.

RETURN VALUES Upon successful completion, the function returns a value of 0. Otherwise the function returns an error number. If the `pthread_kill()` function fails, no signal is sent.

ERRORS The `pthread_kill()` function will fail if:

ESRCH	No thread could be found corresponding to that specified by the given thread ID.
EINVAL	The value of the <i>sig</i> argument is an invalid or unsupported signal number.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `kill(1)`, `pthread_self(3THR)`, `pthread_sigmask(3THR)`, `raise(3C)`, `attributes(5)`, `standards(5)`

pthread_mutexattr_getprioceiling(3THR)

NAME	pthread_mutexattr_getprioceiling, pthread_mutexattr_setprioceiling – get and set prioceiling attribute of mutex attribute object
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_setprioceiling(pthread_mutexattr_t *attr, int prioceiling int *oldceiling); int pthread_mutexattr_getprioceiling(const pthread_mutexattr_t *attr, int *prioceiling);</pre>
DESCRIPTION	<p>The pthread_mutexattr_getprioceiling() and pthread_mutexattr_setprioceiling() functions, respectively, get and set the priority ceiling attribute of a mutex attribute object pointed to by <i>attr</i>, which was previously created by the pthread_mutexattr_init() function.</p> <p>The <i>prioceiling</i> attribute contains the priority ceiling of initialized mutexes. The values of <i>prioceiling</i> must be within the maximum range of priorities defined by SCHED_FIFO.</p> <p>The <i>prioceiling</i> attribute defines the priority ceiling of initialized mutexes, which is the minimum priority level at which the critical section guarded by the mutex is executed. In order to avoid priority inversion, the priority ceiling of the mutex must be set to a priority higher than or equal to the highest priority of all the threads that may lock that mutex. The values of <i>prioceiling</i> must be within the maximum range of priorities defined under the SCHED_FIFO scheduling policy.</p> <p>The ceiling value should be drawn from the range of priorities for the SCHED_FIFO policy. When a thread acquires such a mutex, the policy of the thread at mutex acquisition should match that from which the ceiling value was derived (SCHED_FIFO, in this case). If a thread changes its scheduling policy while holding a ceiling mutex, the behavior of pthread_mutex_lock() and pthread_mutex_unlock() on this mutex is undefined. See pthread_mutex_lock(3THR).</p> <p>The ceiling value should not be treated as a persistent value resident in a pthread_mutex_t that is valid across upgrades of Solaris. The semantics of the actual ceiling value are determined by the existing priority range for the SCHED_FIFO policy, as returned by the sched_get_priority_min() and sched_get_priority_max() functions (see sched_get_priority_min(3RT)) when called on the version of Solaris on which the ceiling value is being utilized.</p>
RETURN VALUES	Upon successful completion, the pthread_mutexattr_getprioceiling() and pthread_mutexattr_setprioceiling() functions return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	The pthread_mutexattr_getprioceiling() and pthread_mutexattr_setprioceiling() functions will fail if:

pthread_mutexattr_getprioceiling(3THR)

ENOSYS The `_POSIX_THREAD_PRIO_PROTECT` option is not defined and the system does not support the function. Since `_POSIX_THREAD_PRIO_PROTECT` is defined, this condition is not reported.

The `pthread_mutexattr_getprioceiling()` function will fail if:

EINVAL The value specified by *attr* is NULL.

The `pthread_mutexattr_setprioceiling()` functions will fail if:

EINVAL The value specified by *attr* is NULL or *prioceiling* is invalid.

The `pthread_mutexattr_getprioceiling()` and `pthread_mutexattr_setprioceiling()` functions may fail if:

EINVAL The value specified by *attr* or *prioceiling* is invalid. This condition is not reported.

EPERM The caller does not have the privilege to perform the operation. This condition is not reported.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `pthread_cond_init(3THR)`, `pthread_create(3THR)`,
`pthread_mutex_init(3THR)`, `pthread_mutex_lock(3THR)`,
`sched_get_priority_min(3RT)`, `attributes(5)`, `standards(5)`

pthread_mutexattr_getprotocol(3THR)

NAME	pthread_mutexattr_getprotocol, pthread_mutexattr_setprotocol – get and set protocol attribute of mutex attribute object
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_getprotocol(const pthread_mutexattr_t *attr, int *protocol); int pthread_mutexattr_setprotocol(pthread_mutexattr_t *attr, int protocol);</pre>
DESCRIPTION	<p>The pthread_mutexattr_setprotocol() and pthread_mutexattr_getprotocol() functions, respectively, set and get the protocol attribute of a mutex attribute object pointed to by <i>attr</i>, which was previously created by the pthread_mutexattr_init() function.</p> <p>The <i>protocol</i> attribute defines the protocol to be followed in utilizing mutexes. The value of <i>protocol</i> may be one of PTHREAD_PRIO_NONE, PTHREAD_PRIO_INHERIT, or PTHREAD_PRIO_PROTECT, which are defined by the header <pthread.h>.</p> <p>When a thread owns a mutex with the PTHREAD_PRIO_NONE protocol attribute, its priority and scheduling are not affected by its mutex ownership.</p> <p>When a thread is blocking higher priority threads because of owning one or more mutexes with the PTHREAD_PRIO_INHERIT protocol attribute, it executes at the higher of its priority or the priority of the highest priority thread waiting on any of the mutexes owned by this thread and initialized with this protocol.</p> <p>When a thread owns one or more mutexes initialized with the PTHREAD_PRIO_PROTECT protocol, it executes at the higher of its priority or the highest of the priority ceilings of all the mutexes owned by this thread and initialized with this attribute, regardless of whether other threads are blocked on any of these mutexes.</p> <p>While a thread is holding a mutex that has been initialized with the PRIORITY_INHERIT or PRIORITY_PROTECT protocol attributes, it will not be subject to being moved to the tail of the scheduling queue at its priority in the event that its original priority is changed, such as by a call to sched_setparam(). Likewise, when a thread unlocks a mutex that has been initialized with the PRIORITY_INHERIT or PRIORITY_PROTECT protocol attributes, it will not be subject to being moved to the tail of the scheduling queue at its priority in the event that its original priority is changed.</p> <p>If a thread simultaneously owns several mutexes initialized with different protocols, it will execute at the highest of the priorities that it would have obtained by each of these protocols.</p> <p>When a thread makes a call to pthread_mutex_lock(), if the symbol _POSIX_THREAD_PRIORITY_INHERIT is defined and the mutex was initialized with the protocol attribute having the value PTHREAD_PRIO_INHERIT, when the calling</p>

pthread_mutexattr_getprotocol(3THR)

thread is blocked because the mutex is owned by another thread, that owner thread will inherit the priority level of the calling thread as long as it continues to own the mutex. The implementation updates its execution priority to the maximum of its assigned priority and all its inherited priorities. Furthermore, if this owner thread becomes blocked on another mutex, the same priority inheritance effect will be propagated to the other owner thread, in a recursive manner.

If the symbol `_POSIX_THREAD_PRIO_INHERIT` is defined, when a mutex initialized with the protocol attribute having the value `PTHREAD_PRIO_INHERIT` dies, the behavior depends on the robustness attribute of the mutex. See `pthread_mutexattr_getrobust_np(3THR)`.

A thread that uses mutexes initialized with the `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT` *protocol* attribute values should have its *contentionscope* attribute equal to `PTHREAD_SCOPE_SYSTEM` (see `pthread_attr_getscope(3THR)`) and its scheduling policy equal to `SCHED_FIFO` or `SCHED_RR` (see `pthread_attr_getschedparam(3THR)` and `pthread_getschedparam(3THR)`).

If a thread with *contentionscope* attribute equal to `PTHREAD_SCOPE_PROCESS` and/or its scheduling policy equal to `SCHED_OTHER` uses a mutex initialized with the `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT` *protocol* attribute value, the effect on the thread's scheduling and priority is unspecified.

The `_POSIX_THREAD_PRIO_INHERIT` and `_POSIX_THREAD_PRIO_PROTECT` options are designed to provide features to solve priority inversion due to mutexes. A priority inheritance or priority ceiling mutex is designed to minimize the dispatch latency of a high priority thread when a low priority thread is holding a mutex required by the high priority thread. This is a specific need for the realtime application domain.

Threads created by realtime applications need to be such that their priorities can influence their access to system resources (CPU resources, at least), in competition with all threads running on the system.

RETURN VALUES

Upon successful completion, the `pthread_mutexattr_getprotocol()` and `pthread_mutexattr_setprotocol()` functions return 0. Otherwise, an error number is returned to indicate the error.

ERRORS

The `pthread_mutexattr_getprotocol()` and `pthread_mutexattr_setprotocol()` functions will fail if:

EINVAL	The value specified by <i>attr</i> is NULL.
ENOSYS	Neither of the options <code>_POSIX_THREAD_PRIO_PROTECT</code> and <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined and the system does not support the function. Since <code>_POSIX_THREAD_PRIO_PROTECT</code> and <code>_POSIX_THREAD_PRIO_INHERIT</code> are defined, this condition is not reported.
ENOTSUP	The value specified by <i>protocol</i> is an unsupported value.

pthread_mutexattr_getprotocol(3THR)

The pthread_mutexattr_getprotocol() and pthread_mutexattr_setprotocol() functions may fail if:

- EINVAL The value specified by *attr* or *protocol* is invalid. This condition is not reported.
- EPERM The caller does not have the privilege to perform the operation. This condition is not reported.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_create(3THR), pthread_mutex_init(3THR), pthread_cond_init(3THR), pthread_mutexattr_getrobust_np(3THR), attributes(5), standards(5)

pthread_mutexattr_getpshared(3THR)

NAME	pthread_mutexattr_getpshared, pthread_mutexattr_setpshared – get and set process-shared attribute
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_mutexattr_getpshared(const pthread_mutexattr_t *attr, int *pshared); int pthread_mutexattr_setpshared(pthread_mutexattr_t *attr, int pshared);</pre>
DESCRIPTION	<p>The <code>pthread_mutexattr_getpshared()</code> function obtains the value of the <i>process-shared</i> attribute from the attributes object referenced by <i>attr</i>. The <code>pthread_mutexattr_setpshared()</code> function is used to set the <i>process-shared</i> attribute in an initialized attributes object referenced by <i>attr</i>.</p> <p>The <i>process-shared</i> attribute is set to <code>PTHREAD_PROCESS_SHARED</code> to permit a mutex to be operated upon by any thread that has access to the memory where the mutex is allocated, even if the mutex is allocated in memory that is shared by multiple processes. If the <i>process-shared</i> attribute is <code>PTHREAD_PROCESS_PRIVATE</code>, the mutex will only be operated upon by threads created within the same process as the thread that initialized the mutex; if threads of differing processes attempt to operate on such a mutex, the behavior is undefined. The default value of the attribute is <code>PTHREAD_PROCESS_PRIVATE</code>.</p>
RETURN VALUES	<p>Upon successful completion, <code>pthread_mutexattr_getpshared()</code> returns 0 and stores the value of the <i>process-shared</i> attribute of <i>attr</i> into the object referenced by the <i>pshared</i> parameter. Otherwise, an error number is returned to indicate the error.</p> <p>Upon successful completion, <code>pthread_mutexattr_setpshared()</code> returns 0. Otherwise, an error number is returned to indicate the error.</p>
ERRORS	<p>The <code>pthread_mutexattr_getpshared()</code> and <code>pthread_mutexattr_setpshared()</code> functions may fail if:</p> <p><code>EINVAL</code> The value specified by <i>attr</i> is invalid.</p> <p>The <code>pthread_mutexattr_setpshared()</code> function may fail if:</p> <p><code>EINVAL</code> The new value specified for the attribute is outside the range of legal values for that attribute.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

`pthread_mutexattr_getpshared(3THR)`

SEE ALSO `pthread_create(3THR)`, `pthread_mutex_init(3THR)`,
`pthread_mutexattr_init(3THR)`, `pthread_cond_init(3THR)`, `attributes(5)`,
`standards(5)`

pthread_mutexattr_getrobust_np(3THR)

NAME	pthread_mutexattr_getrobust_np, pthread_mutexattr_setrobust_np – get or set robustness attribute of mutex attribute object
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_mutexattr_getrobust_np(const pthread_mutexattr_t *attr, int *robustness); int pthread_mutexattr_setrobust_np(pthread_mutexattr_t *attr, int robustness);</pre>
DESCRIPTION	<p>The following applies only if the symbol <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined, and the mutex attributes object <i>attr</i> should be used only to initialize mutexes that will also be initialized with the protocol attribute having the value <code>PTHREAD_PRIO_INHERIT</code>. See <code>pthread_mutexattr_getprotocol(3THR)</code>.</p> <p>The <code>pthread_mutexattr_setrobust_np()</code> and <code>pthread_mutexattr_getrobust_np()</code> functions set and get the <i>robustness</i> attribute of a mutex attribute object pointed to by <i>attr</i> that was previously created by the function <code>pthread_mutexattr_init(3THR)</code>.</p> <p>The <i>robustness</i> attribute defines the behavior when the owner of a mutex dies. The value of <i>robustness</i> may be either <code>PTHREAD_MUTEX_ROBUST_NP</code> or <code>PTHREAD_MUTEX_STALLED_NP</code>, which are defined by the header <code><pthread.h></code>. The default value of the <i>robustness</i> attribute is <code>PTHREAD_MUTEX_STALLED_NP</code>.</p> <p>When the owner of a mutex with the <code>PTHREAD_MUTEX_STALLED_NP</code> <i>robustness</i> attribute dies, all future calls to <code>pthread_mutex_lock(3THR)</code> for this mutex will be blocked from progress in an unspecified manner.</p> <p>When the owner of a mutex with the <code>PTHREAD_MUTEX_ROBUST_NP</code> <i>robustness</i> attribute dies, the mutex is unlocked. The next owner of this mutex acquires it with an error value of <code>EOWNERDEAD</code>. Note that the application must always check the return value from <code>pthread_mutex_lock()</code> for a mutex initialized with the <code>PTHREAD_MUTEX_ROBUST_NP</code> <i>robustness</i> attribute. The new owner of this mutex should then attempt to make the state protected by the mutex consistent, since this state could have been left inconsistent when the last owner died. If the new owner is able to make the state consistent, it should call <code>pthread_mutex_consistent_np(3THR)</code> for the mutex and then unlock the mutex. If for any reason the new owner is not able to make the state consistent, it should not call <code>pthread_mutex_consistent_np()</code> for the mutex, but should simply unlock the mutex. In the latter scenario, all waiters will be awakened and all subsequent calls to <code>pthread_mutex_lock()</code> will fail in acquiring the mutex with an error value of <code>ENOTRECOVERABLE</code>. The mutex can then be made consistent by uninitialized the mutex with the <code>pthread_mutex_destroy()</code> function and reinitializing it with the <code>pthread_mutex_init()</code> function. If the thread that acquired the lock with <code>EOWNERDEAD</code> dies, the next owner will acquire the lock with an error value of <code>EOWNERDEAD</code>.</p>

pthread_mutexattr_getrobust_np(3THR)

Note that the mutex may be in memory shared between processes or in memory private to a process, i.e. the "owner" referenced above is a thread, either within or outside the requestor's process.

The mutex memory must be zeroed before initialization.

RETURN VALUES Upon successful completion, the `pthread_mutexattr_getrobust_np()` and `pthread_mutexattr_setrobust_np()` functions return 0. Otherwise, an error number is returned to indicate the error.

ERRORS The `pthread_mutexattr_getrobust_np()` and `pthread_mutexattr_setrobust_np()` functions will fail if:

- EINVAL** The value specified by *attr* or *robustness* is invalid.
- ENOSYS** The option `_POSIX_THREAD_PRIO_INHERIT` is not defined and the implementation does not support the function.
- ENOTSUP** The value specified by *robustness* is an unsupported value.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `mutex(3THR)`, `pthread_mutex_lock(3THR)`, `pthread_mutex_consistent_np(3THR)`, `pthread_mutexattr_getprotocol(3THR)`, `attributes(5)`, `standards(5)`

pthread_mutexattr_gettype(3THR)

NAME	pthread_mutexattr_gettype, pthread_mutexattr_settype – get or set a mutex type						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_gettype(pthread_mutexattr_t *attr, int *type); int pthread_mutexattr_settype(pthread_mutexattr_t *attr, int type);</pre>						
DESCRIPTION	<p>The pthread_mutexattr_gettype() and pthread_mutexattr_settype() functions respectively get and set the mutex <i>type</i> attribute. This attribute is set in the <i>type</i> parameter to these functions. The default value of the <i>type</i> attribute is PTHREAD_MUTEX_DEFAULT.</p> <p>The type of mutex is contained in the <i>type</i> attribute of the mutex attributes. Valid mutex types include:</p> <table><tr><td>PTHREAD_MUTEX_NORMAL</td><td>This type of mutex does not detect deadlock. A thread attempting to relock this mutex without first unlocking it will deadlock. Attempting to unlock a mutex locked by a different thread results in undefined behavior. Attempting to unlock an unlocked mutex results in undefined behavior.</td></tr><tr><td>PTHREAD_MUTEX_ERRORCHECK</td><td>This type of mutex provides error checking. A thread attempting to relock this mutex without first unlocking it will return with an error. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.</td></tr><tr><td>PTHREAD_MUTEX_RECURSIVE</td><td>A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex. The relocking deadlock that can occur with mutexes of type PTHREAD_MUTEX_NORMAL cannot occur with this type of mutex. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error. This type of mutex is only supported for mutexes whose</td></tr></table>	PTHREAD_MUTEX_NORMAL	This type of mutex does not detect deadlock. A thread attempting to relock this mutex without first unlocking it will deadlock. Attempting to unlock a mutex locked by a different thread results in undefined behavior. Attempting to unlock an unlocked mutex results in undefined behavior.	PTHREAD_MUTEX_ERRORCHECK	This type of mutex provides error checking. A thread attempting to relock this mutex without first unlocking it will return with an error. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.	PTHREAD_MUTEX_RECURSIVE	A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex. The relocking deadlock that can occur with mutexes of type PTHREAD_MUTEX_NORMAL cannot occur with this type of mutex. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error. This type of mutex is only supported for mutexes whose
PTHREAD_MUTEX_NORMAL	This type of mutex does not detect deadlock. A thread attempting to relock this mutex without first unlocking it will deadlock. Attempting to unlock a mutex locked by a different thread results in undefined behavior. Attempting to unlock an unlocked mutex results in undefined behavior.						
PTHREAD_MUTEX_ERRORCHECK	This type of mutex provides error checking. A thread attempting to relock this mutex without first unlocking it will return with an error. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.						
PTHREAD_MUTEX_RECURSIVE	A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex. The relocking deadlock that can occur with mutexes of type PTHREAD_MUTEX_NORMAL cannot occur with this type of mutex. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error. This type of mutex is only supported for mutexes whose						

pthread_mutexattr_gettype(3THR)

process shared attribute is
PTHREAD_PROCESS_PRIVATE.

PTHREAD_MUTEX_DEFAULT

Attempting to recursively lock a mutex of this type results in undefined behavior. Attempting to unlock a mutex of this type that was not locked by the calling thread results in undefined behavior. Attempting to unlock a mutex of this type that is not locked results in undefined behavior. An implementation is allowed to map this mutex to one of the other mutex types.

RETURN VALUES

Upon successful completion, the pthread_mutexattr_settype() function returns 0. Otherwise, an error number is returned to indicate the error.

Upon successful completion, the pthread_mutexattr_gettype() function returns 0 and stores the value of the *type* attribute of *attr* in the object referenced by the *type* parameter. Otherwise an error number is returned to indicate the error.

ERRORS

The pthread_mutexattr_gettype() and pthread_mutexattr_settype() functions will fail if:

EINVAL The value type is invalid.

The pthread_mutexattr_gettype() and pthread_mutexattr_settype() functions may fail if:

EINVAL The value specified by *attr* is invalid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

pthread_cond_timedwait(3THR), pthread_cond_wait(3THR), attributes(5)

NOTES

It is advised that an application should not use a PTHREAD_MUTEX_RECURSIVE mutex with condition variables PTHREAD_MUTEX_RECURSIVE because the implicit unlock performed for a pthread_cond_wait() or pthread_cond_timedwait() will not actually release the mutex (if it had been locked multiple times). If this occurs, no other thread can satisfy the condition of the predicate.

pthread_mutexattr_init(3THR)

NAME	pthread_mutexattr_init, pthread_mutexattr_destroy – initialize and destroy mutex attributes object				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_init(pthread_mutexattr_t *attr); int pthread_mutexattr_destroy(pthread_mutexattr_t *attr);</pre>				
DESCRIPTION	<p>The <code>pthread_mutexattr_init()</code> function initializes a mutex attributes object <i>attr</i> with the default value for all of the attributes defined by the implementation.</p> <p>The effect of initializing an already initialized mutex attributes object is undefined.</p> <p>After a mutex attributes object has been used to initialize one or more mutexes, any function affecting the attributes object (including destruction) does not affect any previously initialized mutexes.</p> <p>The <code>pthread_mutexattr_destroy()</code> function destroys a mutex attributes object; the object becomes, in effect, uninitialized. An implementation may cause <code>pthread_mutexattr_destroy()</code> to set the object referenced by <i>attr</i> to an invalid value. A destroyed mutex attributes object can be re-initialized using <code>pthread_mutexattr_init()</code>; the results of otherwise referencing the object after it has been destroyed are undefined.</p>				
RETURN VALUES	Upon successful completion, <code>pthread_mutexattr_init()</code> and <code>pthread_mutexattr_destroy()</code> return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	The <code>pthread_mutexattr_init()</code> function may fail if: ENOMEM Insufficient memory exists to initialize the mutex attributes object. The <code>pthread_mutexattr_destroy()</code> function may fail if: EINVAL The value specified by <i>attr</i> is invalid.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>pthread_create(3THR)</code> , <code>pthread_mutex_init(3THR)</code> , <code>pthread_mutexattr_init(3THR)</code> , <code>pthread_cond_init(3THR)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

pthread_mutex_consistent_np(3THR)

NAME	pthread_mutex_consistent_np – make a mutex consistent after owner death
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutex_consistent_np(pthread_mutex_t *mutex);</pre>
DESCRIPTION	<p>The following applies only if the symbol <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined, and for mutexes that have been initialized with the protocol attribute having the value <code>PTHREAD_PRIO_INHERIT</code>. See <code>pthread_mutexattr_getprotocol(3THR)</code>.</p> <p>The mutex object referenced by <i>mutex</i> is made consistent by calling <code>pthread_mutex_consistent_np()</code>.</p> <p>A consistent mutex becomes inconsistent and is unlocked if its owner dies while holding it. A subsequent owner of the mutex will acquire the mutex with <code>pthread_mutex_lock(3THR)</code>, which will return <code>EOWNERDEAD</code> to indicate that the acquired mutex is inconsistent.</p> <p>The <code>pthread_mutex_consistent_np()</code> function should be called while holding the mutex acquired by a previous call to <code>pthread_mutex_lock()</code> that returned <code>EOWNERDEAD</code>.</p> <p>Since the critical section protected by the mutex could have been left in an inconsistent state by the dead owner, the caller should make the mutex consistent only if it is able to make the critical section protected by the mutex consistent.</p> <p>Calls to <code>pthread_mutex_lock()</code>, <code>pthread_mutex_unlock()</code>, and <code>pthread_mutex_trylock()</code> for a consistent mutex will behave in the normal manner.</p> <p>The behavior of <code>pthread_mutex_consistent_np()</code> for a mutex which is not inconsistent, or which is not held, is undefined.</p>
RETURN VALUES	Upon successful completion, the <code>pthread_mutexattr_consistent_np()</code> function returns 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The <code>pthread_mutex_consistent_np()</code> function will fail if:</p> <p><code>ENOSYS</code> The option <code>_POSIX_THREAD_PRIO_INHERIT</code> is not defined and the implementation does not support the function.</p> <p>The <code>pthread_mutex_consistent_np()</code> function may fail if:</p> <p><code>EINVAL</code> The value specified by <i>mutex</i> is invalid, or the mutex does not have the appropriate attributes.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

pthread_mutex_consistent_np(3THR)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO mutex(3THR), pthread_mutex_lock(3THR),
pthread_mutexattr_getprotocol(3THR),
pthread_mutexattr_getrobust_np(3THR), attributes(5), standards(5)

pthread_mutex_getprioceiling(3THR)

NAME	pthread_mutex_getprioceiling, pthread_mutex_setprioceiling – change the priority ceiling of a mutex
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutex_setprioceiling(pthread_mutex_t *mutex, int prioceiling, int *old_ceiling); int pthread_mutex_getprioceiling(const pthread_mutex_t *mutex, int *prioceiling);</pre>
DESCRIPTION	<p>The pthread_mutex_getprioceiling() function returns the current priority ceiling of the mutex.</p> <p>The pthread_mutex_setprioceiling() function either locks the mutex if it is unlocked, or blocks until it can successfully lock the mutex, then it changes the mutex's priority ceiling and releases the mutex. When the change is successful, the previous value of the priority ceiling is returned in <i>old_ceiling</i>. The process of locking the mutex need not adhere to the priority protect protocol.</p> <p>If the pthread_mutex_setprioceiling() function fails, the mutex priority ceiling is not changed.</p> <p>The ceiling value should be drawn from the range of priorities for the SCHED_FIFO policy. When a thread acquires such a mutex, the policy of the thread at mutex acquisition should match that from which the ceiling value was derived (SCHED_FIFO, in this case). If a thread changes its scheduling policy while holding a ceiling mutex, the behavior of pthread_mutex_lock() and pthread_mutex_unlock() on this mutex is undefined. See pthread_mutex_lock(3THR).</p> <p>The ceiling value should not be treated as a persistent value resident in a pthread_mutex_t that is valid across upgrades of Solaris. The semantics of the actual ceiling value are determined by the existing priority range for the SCHED_FIFO policy, as returned by the sched_get_priority_min() and sched_get_priority_max() functions (see sched_get_priority_min(3RT)) when called on the version of Solaris on which the ceiling value is being utilized.</p>
RETURN VALUES	Upon successful completion, the pthread_mutex_getprioceiling() and pthread_mutex_setprioceiling() functions return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The pthread_mutex_getprioceiling() and pthread_mutex_setprioceiling() functions will fail if:</p> <p>ENOSYS The option _POSIX_THREAD_PRIO_PROTECT is not defined and the system does not support the function. Since _POSIX_THREAD_PRIO_PROTECT is defined, this condition is not reported.</p>

pthread_mutex_getprioceiling(3THR)

The `pthread_mutex_setprioceiling()` function will fail if:

EINVAL The mutex was not initialized with its *protocol* attribute having the value of `PTHREAD_PRIO_PROTECT`.

The `pthread_mutex_getprioceiling()` and `pthread_mutex_setprioceiling()` functions may fail if:

EINVAL The value specified by *mutex* does not refer to a currently existing mutex. This condition is not reported.

ENOSYS The system does not support the priority ceiling protocol for mutexes. This condition is not reported.

EPERM The caller does not have the privilege to perform the operation. This condition is not reported.

The `pthread_mutex_getprioceiling()` function may fail if:

EINVAL The priority requested by *prioceiling* is out of range. This condition is not reported.

The `pthread_mutex_setprioceiling()` function may fail if:

EINVAL The priority requested by *prioceiling* is out of range. This condition is reported.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `pthread_mutex_init(3THR)`, `pthread_mutex_lock(3THR)`, `sched_get_priority_min(3RT)`, `attributes(5)`, `standards(5)`

NAME	pthread_mutex_init, pthread_mutex_destroy – initialize or destroy a mutex										
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutex_init(pthread_mutex_t *<i>mutex</i>, const pthread_mutexattr_t *<i>attr</i>); int pthread_mutex_destroy(pthread_mutex_t *<i>mutex</i>); pthread_mutex_t <i>mutex</i>= PTHREAD_MUTEX_INITIALIZER</pre>										
DESCRIPTION	<p>The <code>pthread_mutex_init()</code> function initializes the mutex referenced by <i>mutex</i> with attributes specified by <i>attr</i>. If <i>attr</i> is NULL, the default mutex attributes are used; the effect is the same as passing the address of a default mutex attributes object. Upon successful initialization, the state of the mutex becomes initialized and unlocked.</p> <p>Attempting to initialize an already initialized mutex results in undefined behavior.</p> <p>The <code>pthread_mutex_destroy()</code> function destroys the mutex object referenced by <i>mutex</i>; the mutex object becomes, in effect, uninitialized. A destroyed mutex object can be re-initialized using <code>pthread_mutex_init()</code>; the results of otherwise referencing the object after it has been destroyed are undefined.</p> <p>It is safe to destroy an initialized mutex that is unlocked. Attempting to destroy a locked mutex results in undefined behavior.</p> <p>In cases where default mutex attributes are appropriate, the macro <code>PTHREAD_MUTEX_INITIALIZER</code> can be used to initialize mutexes that are statically allocated. The effect is equivalent to dynamic initialization by a call to <code>pthread_mutex_init()</code> with parameter <i>attr</i> specified as NULL, except that no error checks are performed.</p>										
RETURN VALUES	If successful, the <code>pthread_mutex_init()</code> and <code>pthread_mutex_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error.										
ERRORS	<p>The <code>pthread_mutex_init()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EAGAIN</td> <td>The system lacked the necessary resources (other than memory) to initialize another mutex.</td> </tr> <tr> <td>EBUSY</td> <td>An attempt was detected to re-initialize the object referenced by <i>mutex</i>, a robust mutex previously initialized but not yet destroyed.</td> </tr> <tr> <td>EINVAL</td> <td>The value specified by <i>attr</i> has not been initialized using <code>pthread_mutexattr_init(3THR)</code>.</td> </tr> <tr> <td>EINVAL</td> <td>An attempt was made to initialize <i>mutex</i> to be robust and not <code>PTHREAD_PRIO_INHERIT</code>.</td> </tr> <tr> <td>EINVAL</td> <td>An attempt was made to initialize <i>mutex</i> to be <code>PTHREAD_PRIO_INHERIT</code> or <code>PTHREAD_PRIO_PROTECT</code> and the</td> </tr> </table>	EAGAIN	The system lacked the necessary resources (other than memory) to initialize another mutex.	EBUSY	An attempt was detected to re-initialize the object referenced by <i>mutex</i> , a robust mutex previously initialized but not yet destroyed.	EINVAL	The value specified by <i>attr</i> has not been initialized using <code>pthread_mutexattr_init(3THR)</code> .	EINVAL	An attempt was made to initialize <i>mutex</i> to be robust and not <code>PTHREAD_PRIO_INHERIT</code> .	EINVAL	An attempt was made to initialize <i>mutex</i> to be <code>PTHREAD_PRIO_INHERIT</code> or <code>PTHREAD_PRIO_PROTECT</code> and the
EAGAIN	The system lacked the necessary resources (other than memory) to initialize another mutex.										
EBUSY	An attempt was detected to re-initialize the object referenced by <i>mutex</i> , a robust mutex previously initialized but not yet destroyed.										
EINVAL	The value specified by <i>attr</i> has not been initialized using <code>pthread_mutexattr_init(3THR)</code> .										
EINVAL	An attempt was made to initialize <i>mutex</i> to be robust and not <code>PTHREAD_PRIO_INHERIT</code> .										
EINVAL	An attempt was made to initialize <i>mutex</i> to be <code>PTHREAD_PRIO_INHERIT</code> or <code>PTHREAD_PRIO_PROTECT</code> and the										

pthread_mutex_init(3THR)

	type of the mutex to be PTHREAD_MUTEX_ERRORCHECK or PTHREAD_MUTEX_RECURSIVE.
EINVAL	An attempt was made to initialize <i>mutex</i> to be PTHREAD_PROCESS_SHARED and the type of the mutex to be PTHREAD_MUTEX_RECURSIVE.
EINVAL	The process-shared attribute associated with <i>attr</i> is invalid.
ENOMEM	Insufficient memory exists to initialize the mutex.
EPERM	The caller does not have the privilege to perform the operation.
The <code>pthread_mutex_init()</code> function may fail if:	
EBUSY	An attempt was detected to re-initialize the object referenced by <i>mutex</i> , a mutex previously initialized but not yet destroyed. This condition is not reported.
EINVAL	The value specified by <i>attr</i> or <i>mutex</i> is invalid. This condition is not reported.
The <code>pthread_mutex_destroy()</code> function may fail if:	
EBUSY	An attempt was detected to destroy the object referenced by <i>mutex</i> while it is locked or referenced (for example, while being used in a <code>pthread_cond_wait(3THR)</code> or <code>pthread_cond_timedwait(3THR)</code>) by another thread. This condition is not reported.
EINVAL	The value specified by <i>mutex</i> is invalid. This condition is not reported.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `mutex(3THR)`, `pthread_cond_timedwait(3THR)`, `pthread_cond_wait(3THR)`, `pthread_mutex_getprioceiling(3THR)`, `pthread_mutex_lock(3THR)`, `pthread_mutex_unlock(3THR)`, `pthread_mutex_setprioceiling(3THR)`, `pthread_mutex_trylock(3THR)`, `pthread_mutexattr_init(3THR)`, `pthread_mutexattr_getpshared(3THR)`, `pthread_mutexattr_setpshared(3THR)` `attributes(5)`, `standards(5)`

NAME	pthread_mutex_lock, pthread_mutex_trylock, pthread_mutex_unlock – lock or unlock a mutex
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutex_lock(pthread_mutex_t *<i>mutex</i>) ; int pthread_mutex_trylock(pthread_mutex_t *<i>mutex</i>) ; int pthread_mutex_unlock(pthread_mutex_t *<i>mutex</i>) ;</pre>
DESCRIPTION	<p>The mutex object referenced by <i>mutex</i> is locked by calling <code>pthread_mutex_lock()</code>. If the mutex is already locked, the calling thread blocks until the mutex becomes available. This operation returns with the mutex object referenced by <i>mutex</i> in the locked state with the calling thread as its owner.</p> <p>If the mutex type is <code>PTHREAD_MUTEX_NORMAL</code>, deadlock detection is not provided. Attempting to relock the mutex causes deadlock. If a thread attempts to unlock a mutex that it has not locked or a mutex that is unlocked, undefined behavior results.</p> <p>If the mutex type is <code>PTHREAD_MUTEX_ERRORCHECK</code>, then error checking is provided. If a thread attempts to relock a mutex that it has already locked, an error will be returned. If a thread attempts to unlock a mutex that it has not locked or a mutex which is unlocked, an error will be returned.</p> <p>If the mutex type is <code>PTHREAD_MUTEX_RECURSIVE</code>, then the mutex maintains the concept of a lock count. When a thread successfully acquires a mutex for the first time, the lock count is set to 1. Every time a thread relocks this mutex, the lock count is incremented by one. Each time the thread unlocks the mutex, the lock count is decremented by one. When the lock count reaches 0, the mutex becomes available for other threads to acquire. If a thread attempts to unlock a mutex that it has not locked or a mutex that is unlocked, an error will be returned.</p> <p>If the mutex type is <code>PTHREAD_MUTEX_DEFAULT</code>, attempting to recursively lock the mutex results in undefined behavior. Attempting to unlock the mutex if it was not locked by the calling thread results in undefined behavior. Attempting to unlock the mutex if it is not locked results in undefined behavior.</p> <p>The <code>pthread_mutex_trylock()</code> function is identical to <code>pthread_mutex_lock()</code> except that if the mutex object referenced by <i>mutex</i> is currently locked (by any thread, including the current thread), the call returns immediately.</p> <p>The <code>pthread_mutex_unlock()</code> function releases the mutex object referenced by <i>mutex</i>. The manner in which a mutex is released is dependent upon the mutex's type attribute. If there are threads blocked on the mutex object referenced by <i>mutex</i> when <code>pthread_mutex_unlock()</code> is called, resulting in the mutex becoming available, the scheduling policy is used to determine which thread will acquire the mutex. (In the case of <code>PTHREAD_MUTEX_RECURSIVE</code> mutexes, the mutex becomes available when the count reaches 0 and the calling thread no longer has any locks on this mutex.)</p>

pthread_mutex_lock(3THR)

	<p>If a signal is delivered to a thread waiting for a mutex, upon return from the signal handler the thread resumes waiting for the mutex as if it was not interrupted.</p>
RETURN VALUES	<p>If successful, the <code>pthread_mutex_lock()</code> and <code>pthread_mutex_unlock()</code> functions return 0. Otherwise, an error number is returned to indicate the error.</p> <p>The <code>pthread_mutex_trylock()</code> function returns 0 if a lock on the mutex object referenced by <i>mutex</i> is acquired. Otherwise, an error number is returned to indicate the error.</p>
ERRORS	<p>The <code>pthread_mutex_lock()</code> and <code>pthread_mutex_trylock()</code> functions will fail if:</p> <p>EINVAL The <i>mutex</i> was created with the protocol attribute having the value <code>PTHREAD_PRIO_PROTECT</code> and the calling thread's priority is higher than the mutex's current priority ceiling.</p> <p>The <code>pthread_mutex_trylock()</code> function will fail if:</p> <p>EBUSY The <i>mutex</i> could not be acquired because it was already locked.</p> <p>The <code>pthread_mutex_lock()</code>, <code>pthread_mutex_trylock()</code> and <code>pthread_mutex_unlock()</code> functions may fail if:</p> <p>EINVAL The value specified by <i>mutex</i> does not refer to an initialized mutex object.</p> <p>EAGAIN The mutex could not be acquired because the maximum number of recursive locks for <i>mutex</i> has been exceeded.</p> <p>The <code>pthread_mutex_lock()</code> function may fail if:</p> <p>EDEADLK The current thread already owns the mutex.</p> <p>The <code>pthread_mutex_unlock()</code> function may fail if:</p> <p>EPERM The current thread does not own the mutex.</p> <p>When a thread makes a call to <code>pthread_mutex_lock()</code> or <code>pthread_mutex_trylock()</code>, if the symbol <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined and the mutex is initialized with the protocol attribute having the value <code>PTHREAD_PRIO_INHERIT</code> and the robustness attribute having the value <code>PTHREAD_MUTEX_ROBUST_NP</code> (see <code>pthread_mutexattr_getrobust_np(3THR)</code>), the <code>pthread_mutex_lock()</code> and <code>pthread_mutex_trylock()</code> functions will fail if:</p> <p>EOWNERDEAD The last owner of this mutex died while holding the mutex. This mutex is now owned by the caller. The caller must now attempt to make the state protected by the mutex consistent. If it is able to clean up the state, then it should call <code>pthread_mutex_consistent_np()</code> for the mutex</p>

pthread_mutex_lock(3THR)

and unlock the mutex. Subsequent calls to pthread_mutex_lock() and pthread_mutex_trylock() will behave normally, as before. If the caller is not able to clean up the state, pthread_mutex_consistent_np() should not be called for the mutex, but it should be unlocked. Subsequent calls to pthread_mutex_lock() and pthread_mutex_trylock() will fail to acquire the mutex with the error value ENOTRECOVERABLE. If the owner who acquired the lock with EOWNERDEAD dies, the next owner will acquire the lock with EOWNERDEAD.

ENOTRECOVERABLE

The mutex trying to be acquired is protecting the state that has been left irrecoverable by the mutex's last owner, who died while holding the lock. The mutex has not been acquired. This condition can occur when the lock was previously acquired with EOWNERDEAD, and the owner was not able to clean up the state and unlocked the mutex without making the mutex consistent.

ENOMEM

The limit on the number of simultaneously held mutexes has been exceeded.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_mutex_init(3THR), pthread_mutex_destroy(3THR), pthread_mutex_consistent_np(3THR), pthread_mutexattr_getrobust_np(3THR), attributes(5), standards(5)

NOTES In the current implementation of threads, pthread_mutex_lock(), pthread_mutex_unlock(), mutex_lock(), mutex_unlock(), pthread_mutex_trylock(), and mutex_trylock() do not validate the mutex type. Therefore, an uninitialized mutex or a mutex with an invalid type does not return EINVAL. Interfaces for mutexes with an invalid type have unspecified behavior.

Uninitialized mutexes that are allocated locally may contain junk data. Such mutexes need to be initialized using pthread_mutex_init() or mutex_init().

pthread_once(3THR)

NAME	pthread_once – initialize dynamic package				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> pthread_once_t <i>once_control</i> = PTHREAD_ONCE_INIT; int pthread_once(pthread_once_t *<i>once_control</i>, void (*<i>init_routine</i>, void));</pre>				
DESCRIPTION	<p>If any thread in a process with a <i>once_control</i> parameter makes a call to <code>pthread_once()</code>, the first call will summon the <code>init_routine()</code>, but subsequent calls will not. The <i>once_control</i> parameter determines whether the associated initialization routine has been called. The <code>init_routine()</code> is complete upon return of <code>pthread_once()</code>.</p> <p><code>pthread_once()</code> is not a cancellation point; however, if the function <code>init_routine()</code> is a cancellation point and is canceled, the effect on <i>once_control</i> is the same as if <code>pthread_once()</code> had never been called.</p> <p>The constant <code>PTHREAD_ONCE_INIT</code> is defined in the <code><pthread.h></code> header.</p> <p>If <i>once_control</i> has automatic storage duration or is not initialized by <code>PTHREAD_ONCE_INIT</code>, the behavior of <code>pthread_once()</code> is undefined.</p>				
RETURN VALUES	Upon successful completion, <code>pthread_once()</code> returns 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<code>EINVAL</code> <i>once_control</i> or <i>init_routine</i> is NULL.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>attributes(5)</code>				
NOTES	Solaris threads do not offer this functionality.				

pthread_rwlockattr_getpshared(3THR)

NAME pthread_rwlockattr_getpshared, pthread_rwlockattr_setpshared – get or set process-shared attribute of read-write lock attributes object

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

int pthread_rwlockattr_getpshared(const pthread_rwlockattr_t
    *attr, int *pshared);

int pthread_rwlockattr_setpshared(pthread_rwlockattr_t *attr, int
    pshared);
```

DESCRIPTION

The *process-shared* attribute is set to PTHREAD_PROCESS_SHARED to permit a read-write lock to be operated upon by any thread that has access to the memory where the read-write lock is allocated, even if the read-write lock is allocated in memory that is shared by multiple processes. If the *process-shared* attribute is PTHREAD_PROCESS_PRIVATE, the read-write lock will only be operated upon by threads created within the same process as the thread that initialised the read-write lock; if threads of differing processes attempt to operate on such a read-write lock, the behaviour is undefined. The default value of the *process-shared* attribute is PTHREAD_PROCESS_PRIVATE.

The pthread_rwlockattr_getpshared() function obtains the value of the *process-shared* attribute from the initialised attributes object referenced by *attr*. The pthread_rwlockattr_setpshared() function is used to set the *process-shared* attribute in an initialised attributes object referenced by *attr*.

RETURN VALUES

If successful, the pthread_rwlockattr_setpshared() function returns 0. Otherwise, an error number is returned to indicate the error.

Upon successful completion, the pthread_rwlockattr_getpshared() returns 0 and stores the value of the *process-shared* attribute of *attr* into the object referenced by the *pshared* parameter. Otherwise an error number is returned to indicate the error.

ERRORS

The pthread_rwlockattr_getpshared() and pthread_rwlockattr_setpshared() functions will fail if:

EINVAL The value specified by *attr* or *pshared* is invalid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

pthread_rwlock_init(3THR), pthread_rwlock_rdlock(3THR), pthread_rwlock_unlock(3THR), pthread_rwlock_wrlock(3THR), pthread_rwlockattr_init(3THR), attributes(5)

pthread_rwlockattr_init(3THR)

NAME	pthread_rwlockattr_init, pthread_rwlockattr_destroy – initialize or destroy read-write lock attributes object				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_rwlockattr_init(pthread_rwlockattr_t *<i>attr</i>); int pthread_rwlockattr_destroy(pthread_rwlockattr_t *<i>attr</i>);</pre>				
DESCRIPTION	<p>The <code>pthread_rwlockattr_init()</code> function initializes a read-write lock attributes object <i>attr</i> with the default value for all of the attributes defined by the implementation.</p> <p>Results are undefined if <code>pthread_rwlockattr_init()</code> is called specifying an already initialized read-write lock attributes object.</p> <p>After a read-write lock attributes object has been used to initialize one or more read-write locks, any function affecting the attributes object (including destruction) does not affect any previously initialized read-write locks.</p> <p>The <code>pthread_rwlockattr_destroy()</code> function destroys a read-write lock attributes object. The effect of subsequent use of the object is undefined until the object is re-initialized by another call to <code>pthread_rwlockattr_init()</code>. An implementation may cause <code>pthread_rwlockattr_destroy()</code> to set the object referenced by <i>attr</i> to an invalid value.</p>				
RETURN VALUES	If successful, the <code>pthread_rwlockattr_init()</code> and <code>pthread_rwlockattr_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The <code>pthread_rwlockattr_init()</code> function will fail if:</p> <table><tr><td>ENOMEM</td><td>Insufficient memory exists to initialize the read-write lock attributes object.</td></tr></table> <p>The <code>pthread_rwlockattr_destroy()</code> function may fail if:</p> <table><tr><td>EINVAL</td><td>The value specified by <i>attr</i> is invalid.</td></tr></table>	ENOMEM	Insufficient memory exists to initialize the read-write lock attributes object.	EINVAL	The value specified by <i>attr</i> is invalid.
ENOMEM	Insufficient memory exists to initialize the read-write lock attributes object.				
EINVAL	The value specified by <i>attr</i> is invalid.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>pthread_rwlock_init(3THR)</code> , <code>pthread_rwlock_rdlock(3THR)</code> , <code>pthread_rwlock_unlock(3THR)</code> , <code>pthread_rwlock_wrlock(3THR)</code> , <code>pthread_rwlockattr_getpshared(3THR)</code> , <code>attributes(5)</code>				

NAME	pthread_rwlock_init, pthread_rwlock_destroy – initialize or destroy a read-write lock object				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_rwlock_init(pthread_rwlock_t *<i>rwlock</i>, const pthread_rwlockattr_t *<i>attr</i>); int pthread_rwlock_destroy(pthread_rwlock_t *<i>rwlock</i>); pthread_rwlock_t <i>rwlock</i>=PTHREAD_RWLOCK_INITIALIZER;</pre>				
DESCRIPTION	<p>The <code>pthread_rwlock_init()</code> function initializes the read-write lock referenced by <i>rwlock</i> with the attributes referenced by <i>attr</i>. If <i>attr</i> is NULL, the default read-write lock attributes are used; the effect is the same as passing the address of a default read-write lock attributes object. Once initialized, the lock can be used any number of times without being re-initialized. Upon successful initialization, the state of the read-write lock becomes initialized and unlocked. Results are undefined if <code>pthread_rwlock_init()</code> is called specifying an already initialized read-write lock. Results are undefined if a read-write lock is used without first being initialized.</p> <p>If the <code>pthread_rwlock_init()</code> function fails, <i>rwlock</i> is not initialized and the contents of <i>rwlock</i> are undefined.</p> <p>The <code>pthread_rwlock_destroy()</code> function destroys the read-write lock object referenced by <i>rwlock</i> and releases any resources used by the lock. The effect of subsequent use of the lock is undefined until the lock is re-initialized by another call to <code>pthread_rwlock_init()</code>. An implementation may cause <code>pthread_rwlock_destroy()</code> to set the object referenced by <i>rwlock</i> to an invalid value. Results are undefined if <code>pthread_rwlock_destroy()</code> is called when any thread holds <i>rwlock</i>. Attempting to destroy an uninitialized read-write lock results in undefined behaviour. A destroyed read-write lock object can be re-initialized using <code>pthread_rwlock_init()</code>; the results of otherwise referencing the read-write lock object after it has been destroyed are undefined.</p> <p>In cases where default read-write lock attributes are appropriate, the macro <code>PTHREAD_RWLOCK_INITIALIZER</code> can be used to initialize read-write locks that are statically allocated. The effect is equivalent to dynamic initialization by a call to <code>pthread_rwlock_init()</code> with the parameter <i>attr</i> specified as NULL, except that no error checks are performed.</p>				
RETURN VALUES	If successful, the <code>pthread_rwlock_init()</code> and <code>pthread_rwlock_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The <code>pthread_rwlock_init()</code> and <code>pthread_rwlock_destroy()</code> functions will fail if:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The value specified by <i>attr</i> is invalid.</td> </tr> <tr> <td>EINVAL</td> <td>The value specified by <i>rwlock</i> is invalid.</td> </tr> </table>	EINVAL	The value specified by <i>attr</i> is invalid.	EINVAL	The value specified by <i>rwlock</i> is invalid.
EINVAL	The value specified by <i>attr</i> is invalid.				
EINVAL	The value specified by <i>rwlock</i> is invalid.				

pthread_rwlock_init(3THR)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_rwlock_rdlock(3THR), pthread_rwlock_unlock(3THR), pthread_rwlock_wrlock(3THR), pthread_rwlockattr_init(3THR), attributes(5)

NAME	pthread_rwlock_rdlock, pthread_rwlock_tryrdlock – lock or attempt to lock a read-write lock object for reading
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_rwlock_rdlock(pthread_rwlock_t *<i>rwlock</i>); int pthread_rwlock_tryrdlock(pthread_rwlock_t *<i>rwlock</i>);</pre>
DESCRIPTION	<p>The pthread_rwlock_rdlock() function applies a read lock to the read-write lock referenced by <i>rwlock</i>. The calling thread acquires the read lock if a writer does not hold the lock and there are no writers blocked on the lock. It is unspecified whether the calling thread acquires the lock when a writer does not hold the lock and there are writers waiting for the lock. If a writer holds the lock, the calling thread will not acquire the read lock. If the read lock is not acquired, the calling thread blocks (that is, it does not return from the pthread_rwlock_rdlock() call) until it can acquire the lock. Results are undefined if the calling thread holds a write lock on <i>rwlock</i> at the time the call is made.</p> <p>Implementations are allowed to favor writers over readers to avoid writer starvation. The current implementation favors writers over readers.</p> <p>A thread may hold multiple concurrent read locks on <i>rwlock</i> (that is, successfully call the pthread_rwlock_rdlock() function <i>n</i> times). If so, the thread must perform matching unlocks (that is, it must call the pthread_rwlock_unlock() function <i>n</i> times).</p> <p>The function pthread_rwlock_tryrdlock() applies a read lock as in the pthread_rwlock_rdlock() function with the exception that the function fails if any thread holds a write lock on <i>rwlock</i> or there are writers blocked on <i>rwlock</i>.</p> <p>Results are undefined if any of these functions are called with an uninitialized read-write lock.</p> <p>If a signal is delivered to a thread waiting for a read-write lock for reading, upon return from the signal handler the thread resumes waiting for the read-write lock for reading as if it was not interrupted.</p>
RETURN VALUES	<p>If successful, the pthread_rwlock_rdlock() function returns 0. Otherwise, an error number is returned to indicate the error.</p> <p>The function pthread_rwlock_tryrdlock() returns 0 if the lock for reading on the read-write lock object referenced by <i>rwlock</i> is acquired. Otherwise an error number is returned to indicate the error.</p>
ERRORS	<p>The pthread_rwlock_tryrdlock() function will fail if:</p> <p>EBUSY The read-write lock could not be acquired for reading because a writer holds the lock or was blocked on it.</p>

pthread_rwlock_rdlock(3THR)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO | `pthread_rwlock_init(3THR)`, `pthread_rwlock_wrlock(3THR)`,
`pthread_rwlockattr_init(3THR)`, `pthread_rwlock_unlock(3THR)`,
`attributes(5)`

NAME	pthread_rwlock_unlock – unlock a read-write lock object				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_rwlock_unlock(pthread_rwlock_t *<i>rwlock</i>);</pre>				
DESCRIPTION	<p>The <code>pthread_rwlock_unlock()</code> function is called to release a lock held on the read-write lock object referenced by <i>rwlock</i>. Results are undefined if the read-write lock <i>rwlock</i> is not held by the calling thread.</p> <p>If this function is called to release a read lock from the read-write lock object and there are other read locks currently held on this read-write lock object, the read-write lock object remains in the read locked state. If this function releases the calling thread's last read lock on this read-write lock object, then the calling thread is no longer one of the owners of the object. If this function releases the last read lock for this read-write lock object, the read-write lock object will be put in the unlocked state with no owners.</p> <p>If this function is called to release a write lock for this read-write lock object, the read-write lock object will be put in the unlocked state with no owners.</p> <p>If the call to the <code>pthread_rwlock_unlock()</code> function results in the read-write lock object becoming unlocked and there are multiple threads waiting to acquire the read-write lock object for writing, the scheduling policy is used to determine which thread acquires the read-write lock object for writing. If there are multiple threads waiting to acquire the read-write lock object for reading, the scheduling policy is used to determine the order in which the waiting threads acquire the read-write lock object for reading. If there are multiple threads blocked on <i>rwlock</i> for both read locks and write locks, it is unspecified whether the readers acquire the lock first or whether a writer acquires the lock first.</p> <p>Results are undefined if any of these functions are called with an uninitialized read-write lock.</p>				
RETURN VALUES	If successful, the <code>pthread_rwlock_unlock()</code> function returns 0. Otherwise, an error number is returned to indicate the error.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>pthread_rwlock_init(3THR)</code> , <code>pthread_rwlock_rdlock(3THR)</code> , <code>pthread_rwlock_wrlock(3THR)</code> , <code>pthread_rwlockattr_init(3THR)</code> , <code>attributes(5)</code>				

pthread_rwlock_wrlock(3THR)

NAME	pthread_rwlock_wrlock, pthread_rwlock_trywrlock – lock or attempt to lock a read-write lock object for writing
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_rwlock_wrlock(pthread_rwlock_t *rwlock); int pthread_rwlock_trywrlock(pthread_rwlock_t *rwlock);</pre>
DESCRIPTION	<p>The <code>pthread_rwlock_wrlock()</code> function applies a write lock to the read-write lock referenced by <code>rwlock</code>. The calling thread acquires the write lock if no other thread (reader or writer) holds the read-write lock <code>rwlock</code>. Otherwise, the thread blocks (that is, does not return from the <code>pthread_rwlock_wrlock()</code> call) until it can acquire the lock. Results are undefined if the calling thread holds the read-write lock (whether a read or write lock) at the time the call is made.</p> <p>Implementations are allowed to favor writers over readers to avoid writer starvation. The current implementation favors writers over readers.</p> <p>The function <code>pthread_rwlock_trywrlock()</code> applies a write lock like the <code>pthread_rwlock_wrlock()</code> function, with the exception that the function fails if any thread currently holds <code>rwlock</code> (for reading or writing).</p> <p>Results are undefined if any of these functions are called with an uninitialized read-write lock.</p> <p>If a signal is delivered to a thread waiting for a read-write lock for writing, upon return from the signal handler the thread resumes waiting for the read-write lock for writing as if it was not interrupted.</p>
RETURN VALUES	<p>If successful, the <code>pthread_rwlock_wrlock()</code> function returns 0. Otherwise, an error number is returned to indicate the error.</p> <p>The function <code>pthread_rwlock_trywrlock()</code> returns 0 if the lock for writing on the read-write lock object referenced by <code>rwlock</code> is acquired. Otherwise an error number is returned to indicate the error.</p>
ERRORS	<p>The <code>pthread_rwlock_trywrlock()</code> function will fail if:</p> <p>EBUSY The read-write lock could not be acquired for writing because it was already locked for reading or writing.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

pthread_rwlock_wrlock(3THR)

SEE ALSO pthread_rwlock_init(3THR), pthread_rwlock_unlock(3THR),
pthread_rwlockattr_init(3THR), pthread_rwlock_rdlock(3THR),
attributes(5)

pthread_self(3THR)

NAME pthread_self – get calling thread’s ID

SYNOPSIS `cc -mt [flag...] file... -lpthread [-lrt library...]`

```
#include <pthread.h>
pthread_t pthread_self(void);
```

DESCRIPTION The pthread_self() function returns the thread ID of the calling thread.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_create(3THR), pthread_equal(3THR), attributes(5), standards(5)

NAME	pthread_setcancelstate – enable or disable cancellation
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_setcancelstate(int <i>state</i>, int *<i>oldstate</i>);</pre>
DESCRIPTION	<p>pthread_setcancelstate() atomically sets the calling thread's cancellation state to the specified <i>state</i> and if <i>oldstate</i> is not NULL, stores the previous cancellation <i>state</i> in <i>oldstate</i>.</p> <p>The <i>state</i> can be either of the following:</p> <p>PTHREAD_CANCEL_ENABLE This is the default. When cancellation is deferred (deferred cancellation is also the default), cancellation occurs when the target thread reaches a cancellation point and a cancel is pending. When cancellation is asynchronous, receipt of a pthread_cancel(3THR) call causes immediate cancellation.</p> <p>PTHREAD_CANCEL_DISABLE When cancellation is deferred, all cancellation requests to the target thread are held pending. When cancellation is asynchronous, all cancellation requests to the target thread are held pending; as soon as cancellation is re-enabled, pending cancellations are executed immediately.</p> <p>See cancellation(3THR) for the definition of a cancellation point and a discussion of cancellation concepts. See pthread_setcanceltype(3THR) for explanations of deferred and asynchronous cancellation.</p> <p>pthread_setcancelstate() is a cancellation point when it is called with PTHREAD_CANCEL_ENABLE and the cancellation type is PTHREAD_CANCEL_ASYNCHRONOUS.</p>
RETURN VALUES	Upon successful completion, pthread_setcancelstate(), returns 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The pthread_setcancelstate() function will fail if:</p> <p>EINVAL The specified <i>state</i> is not PTHREAD_CANCEL_ENABLE or PTHREAD_CANCEL_DISABLE.</p>
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

pthread_setcancelstate(3THR)

SEE ALSO | cancellation(3THR), condition(3THR), pthread_cancel(3THR),
pthread_cleanup_pop(3THR), pthread_cleanup_push(3THR),
pthread_exit(3THR), pthread_join(3THR), pthread_setcanceltype(3THR),
pthread_testcancel(3THR), setjmp(3C), attributes(5)

pthread_setcanceltype(3THR)

NAME pthread_setcanceltype – set the cancellation type of a thread

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
```

```
#include <pthread.h>

int pthread_setcanceltype(int type, int *oldtype);
```

DESCRIPTION pthread_setcanceltype() atomically sets the calling thread's cancellation type to the specified *type* and, if *oldtype* is not NULL, stores the previous cancellation type in *oldtype*. The *type* can be either of the following:

PTHREAD_CANCEL_DEFERRED	This is the default. When cancellation is enabled (enabled cancellation is also the default), cancellation occurs when the target thread reaches a cancellation point and a cancel is pending. When cancellation is disabled, all cancellation requests to the target thread are held pending.
PTHREAD_CANCEL_ASYNCHRONOUS	When cancellation is enabled, receipt of a pthread_cancel(3THR) call causes immediate cancellation. When cancellation is disabled, all cancellation requests to the target thread are held pending; as soon as cancellation is re-enabled, pending cancellations are executed immediately.

See cancellation(3THR) for the definition of a cancellation point and a discussion of cancellation concepts. See pthread_setcancelstate(3THR) for explanations of enabling and disabling cancellation.

pthread_setcanceltype() is a cancellation point if *type* is called with PTHREAD_CANCEL_ASYNCHRONOUS and the cancellation state is PTHREAD_CANCEL_ENABLE.

RETURN VALUES Upon successful completion, the pthread_setcanceltype() function returns 0. Otherwise, an error number is returned to indicate the error.

ERRORS The pthread_setcanceltype() function will fail if:

EINVAL	The specified <i>type</i> is not PTHREAD_CANCEL_DEFERRED or PTHREAD_CANCEL_ASYNCHRONOUS.
--------	--

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

pthread_setcanceltype(3THR)

SEE ALSO | cancellation(3THR), condition(3THR), pthread_cancel(3THR),
pthread_cleanup_pop(3THR), pthread_cleanup_push(3THR),
pthread_exit(3THR), pthread_join(3THR),
pthread_setcancelstate(3THR), pthread_testcancel(3THR), setjmp(3C),
attributes(5)

NAME	pthread_sigmask – change or examine calling thread’s signal mask						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> #include <signal.h> int pthread_sigmask(int <i>how</i>, const sigset_t *<i>set</i>, sigset_t *<i>oset</i>);</pre>						
DESCRIPTION	<p>The <code>pthread_sigmask()</code> function changes or examines a calling thread’s signal mask. Each thread has its own signal mask. A new thread inherits the calling thread’s signal mask and priority; however, pending signals are not inherited. Signals pending for a new thread will be empty.</p> <p>If the value of the argument <code>set</code> is not <code>NULL</code>, <code>set</code> points to a set of signals that can modify the currently blocked set. If the value of <code>set</code> is <code>NULL</code>, the value of <code>how</code> is insignificant and the thread’s signal mask is unmodified; thus, <code>pthread_sigmask()</code> can be used to inquire about the currently blocked signals.</p> <p>The value of the argument <code>how</code> specifies the method in which the set is changed and takes one of the following values:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">SIG_BLOCK</td> <td>set corresponds to a set of signals to block. They are added to the current signal mask.</td> </tr> <tr> <td>SIG_UNBLOCK</td> <td>set corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.</td> </tr> <tr> <td>SIG_SETMASK</td> <td>set corresponds to the new signal mask. The current signal mask is replaced by <code>set</code>.</td> </tr> </table> <p>If the value of <code>oset</code> is not <code>NULL</code>, it points to the location where the previous signal mask is stored.</p>	SIG_BLOCK	set corresponds to a set of signals to block. They are added to the current signal mask.	SIG_UNBLOCK	set corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.	SIG_SETMASK	set corresponds to the new signal mask. The current signal mask is replaced by <code>set</code> .
SIG_BLOCK	set corresponds to a set of signals to block. They are added to the current signal mask.						
SIG_UNBLOCK	set corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.						
SIG_SETMASK	set corresponds to the new signal mask. The current signal mask is replaced by <code>set</code> .						
RETURN VALUES	Upon successful completion, the <code>pthread_sigmask()</code> function returns 0. Otherwise, it returns a non-zero value.						
ERRORS	<p>The <code>pthread_sigmask()</code> function will fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The value of <code>how</code> is not defined and <code>oset</code> is <code>NULL</code>.</td> </tr> </table>	EINVAL	The value of <code>how</code> is not defined and <code>oset</code> is <code>NULL</code> .				
EINVAL	The value of <code>how</code> is not defined and <code>oset</code> is <code>NULL</code> .						
EXAMPLES	<p>EXAMPLE 1 The following example shows how to create a default thread that can serve as a signal catcher/handler with its own signal mask. <code>new</code> will have a different value from the creator’s signal mask.</p> <p>As POSIX threads and Solaris threads are fully compatible even within the same process, this example uses <code>pthread_create(3THR)</code> if you execute <code>a.out 0</code>, or <code>thr_create(3THR)</code> if you execute <code>a.out 1</code>.</p> <p>In this example:</p> <ul style="list-style-type: none"> ■ <code>sigemptyset(3C)</code> initializes a null signal set, <code>new</code>. <code>sigaddset(3C)</code> packs the signal, <code>SIGINT</code>, into that new set. 						

pthread_sigmask(3THR)

EXAMPLE 1 The following example shows how to create a default thread that can serve as a signal catcher/handler with its own signal mask. `new` will have a different value from the creator's signal mask. (Continued)

- Either `pthread_sigmask()` or `thr_sigsetmask()` is used to mask the signal, `SIGINT` (CTRL-C), from the calling thread, which is `main()`. The signal is masked to guarantee that only the new thread will receive this signal.
- `pthread_create()` or `thr_create()` creates the signal-handling thread.
- Using `pthread_join(3THR)` or `thr_join(3THR)`, `main()` then waits for the termination of that signal-handling thread, whose ID number is `user_threadID`; after which, `main()` will `sleep(3C)` for 2 seconds, and then the program terminates.
- The signal-handling thread, handler:
 - Assigns the handler `interrupt()` to handle the signal `SIGINT`, by the call to `sigaction(2)`.
 - Resets its own signal set to *not block* the signal, `SIGINT`.
 - Sleeps for 8 seconds to allow time for the user to deliver the signal, `SIGINT`, by pressing the CTRL-C.

```
/* cc thisfile.c -lthread -lpthread */
#define _REENTRANT /* basic first 3-lines for threads */
#include <pthread.h>
#include <thread.h>
thread_t user_threadID;
sigset_t new;
void *handler( ), interrupt( );

main( int argc, char *argv[ ] ) {
    test_argv(argv[1]);

    sigemptyset(&new);
    sigaddset(&new, SIGINT);
    switch(*argv[1]) {

        case '0': /* POSIX */
            pthread_sigmask(SIG_BLOCK, &new, NULL);
            pthread_create(&user_threadID, NULL, handler, argv[1]);
            pthread_join(user_threadID, NULL);
            break;

        case '1': /* Solaris */
            thr_sigsetmask(SIG_BLOCK, &new, NULL);
            thr_create(NULL, 0, handler, argv[1], 0, &user_threadID);
            thr_join(user_threadID, NULL, NULL);
            break;
    } /* switch */

    printf("thread handler, # %d, has exited\n",user_threadID);
    sleep(2);
    printf("main thread, # %d is done\n", thr_self( ));
} /* end main */
```

EXAMPLE 1 The following example shows how to create a default thread that can serve as a signal catcher/handler with its own signal mask. `new` will have a different value from the creator's signal mask. (Continued)

```

struct sigaction act;

void *
handler(char argv[ ])
{
    act.sa_handler = interrupt;
    sigaction(SIGINT, &act, NULL);
    switch(*argv) {
        case '0': /* POSIX */
            pthread_sigmask(SIG_UNBLOCK, &new, NULL);
            break;
        case '1': /* Solaris */
            thr_sigsetmask(SIG_UNBLOCK, &new, NULL);
            break;
    }
    printf("\n Press CTRL-C to deliver SIGINT signal to the process\n");
    sleep(8); /* give user time to hit CTRL-C */
}

void
interrupt(int sig)
{
    printf("thread %d caught signal %d\n", thr_self( ), sig);
}

void test_argv(char argv[ ]) {
    if(argv == NULL) {
        printf("use 0 as arg1 to use thr_create( );\n \
or use 1 as arg1 to use pthread_create( )\n");
        exit(NULL);
    }
}

```

EXAMPLE 2

In the last example, the `handler` thread served as a signal-handler while also taking care of activity of its own (in this case, sleeping, although it could have been some other activity). A thread could be completely dedicated to signal-handling simply by waiting for the delivery of a selected signal by blocking with `sigwait(2)`. The two subroutines in the previous example, `handler()` and `interrupt()`, could have been replaced with the following routine:

```

void *
handler( )
{
    int signal;
    printf("thread %d is waiting for you to press the CTRL-C keys\n",
           thr_self( ));
    sigwait(&new, &signal);
    printf("thread %d has received the signal %d \n", thr_self( ), signal);
}

```

pthread_sigmask(3THR)

EXAMPLE 2 (Continued)

```
}
/* pthread_create( ) and thr_create( ) would use NULL instead of argv[1]
   for the arg passed to handler( ) */
```

In this routine, one thread is dedicated to catching and handling the signal specified by the set new, which allows main() and all of its other sub-threads, created *after* pthread_sigmask() or thr_sigsetmask() masked that signal, to continue uninterrupted. Any use of sigwait(2) should be such that all threads block the signals passed to sigwait(2) at all times. Only the thread that calls sigwait() will get the signals. The call to sigwait(2) takes two arguments.

For this type of background dedicated signal-handling routine, you may wish to use a Solaris daemon thread by passing the argument, THR_DAEMON, to thr_create(3THR).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe and Async-Signal-Safe

SEE ALSO sigaction(2), sigprocmask(2), sigwait(2), cond_wait(3THR), pthread_create(3THR), pthread_join(3THR), pthread_self(3THR), sigsetops(3C), sleep(3C), attributes(5), standards(5)

NOTES It is not possible to block signals that cannot be ignored (see sigaction(2)). If using the threads library, it is not possible to block or unblock the signals SIGWAITING, SIGLWP, or SIGCANCEL. This restriction is quietly enforced by the threads library.

Using sigwait(2) in a dedicated thread allows asynchronously generated signals to be managed synchronously; however, sigwait(2) should never be used to manage synchronously generated signals.

Synchronously generated signals are exceptions that are generated by a thread and are directed at the thread causing the exception. Since sigwait() blocks waiting for signals, the blocking thread cannot receive a synchronously generated signal.

If sigprocmask(2) is used in a multi-threaded program, it will be the same as if pthread_sigmask() has been called. POSIX leaves the semantics of the call to sigprocmask(2) unspecified in a multi-threaded process, so programs that care about POSIX portability should not depend on this semantic.

If a signal is delivered while a thread is waiting on a condition variable, the cond_wait(3THR) function will be interrupted and the handler will be executed. The state of the lock protecting the condition variable is undefined while the thread is executing the signal handler.

pthread_sigmask(3THR)

Although `pthread_sigmask()` is Async-Signal-Safe with respect to the Solaris environment, this safeness is not guaranteed to be portable to other POSIX domains.

Signals that are generated synchronously should not be masked. If such a signal is blocked and delivered, the receiving process is killed.

pthread_testcancel(3THR)

NAME	pthread_testcancel – create cancellation point in the calling thread				
SYNOPSIS	<pre>cc -mt [flag...] file... -pthread [-lrt library...] #include <pthread.h> void pthread_testcancel (void);</pre>				
DESCRIPTION	The pthread_testcancel() function forces testing for cancellation. This is useful when you need to execute code that runs for long periods without encountering cancellation points; such as a library routine that executes long-running computations without cancellation points. This type of code can block cancellation for unacceptable long periods of time. One strategy for avoiding blocking cancellation for long periods, is to insert calls to pthread_testcancel() in the long-running computation code and to setup a cancellation handler in the library code, if required.				
RETURN VALUES	The pthread_testcancel() function returns a void.				
ERRORS	The pthread_testcancel() function does not return errors.				
EXAMPLES	EXAMPLE 1 See cancellation(3THR) for an example of using pthread_testcancel() to force testing for cancellation and a discussion of cancellation concepts.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	Intro(3), cancellation(3THR), condition(3THR), pthread_cleanup_pop(3THR), pthread_cleanup_push(3THR), pthread_exit(3THR), pthread_join(3THR), pthread_setcancelstate(3THR), pthread_setcanceltype(3THR), setjmp(3C), attributes(5)				
NOTES	<p>pthread_testcancel() has no effect if cancellation is disabled.</p> <p>Use pthread_testcancel() with pthread_setcanceltype() called with its canceltype set to PTHREAD_CANCEL_DEFERRED. pthread_testcancel() operation is undefined if pthread_setcanceltype() was called with its canceltype argument set to PTHREAD_CANCEL_ASYNCHRONOUS.</p> <p>It is possible to kill a thread when it is holding a resource, such as lock or allocated memory. If that thread has not setup a cancellation cleanup handler to release the held resource, the application is "cancel-unsafe". See attributes(5) for a discussion of Cancel-Safety, Deferred-Cancel-Safety, and Asynchronous-Cancel-Safety.</p>				

NAME	rwlock, rwlock_init, rwlock_destroy, rw_rdlock, rw_wrlock, rw_tryrdlock, rw_trywrlock, rw_unlock – multiple readers, single writer locks				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <synch.h> int rwlock_init(rwlock_t *rwlp, int type, void * arg); int rwlock_destroy(rwlock_t *rwlp); int rw_rdlock(rwlock_t *rwlp); int rw_wrlock(rwlock_t *rwlp); int rw_unlock(rwlock_t *rwlp); int rw_tryrdlock(rwlock_t *rwlp); int rw_trywrlock(rwlock_t *rwlp);</pre>				
DESCRIPTION	<p>Many threads can have simultaneous read-only access to data, while only one thread can have write access at any given time. Multiple read access with single write access is controlled by locks, which are generally used to protect data that is frequently searched.</p> <p>Readers/writer locks can synchronize threads in this process and other processes if they are allocated in writable memory and shared among cooperating processes (see <code>mmap(2)</code>), and are initialized for this purpose.</p> <p>Additionally, readers/writer locks must be initialized prior to use. <code>rwlock_init()</code> The readers/writer lock pointed to by <code>rwlp</code> is initialized by <code>rwlock_init()</code>. A readers/writer lock is capable of having several types of behavior, which is specified by <code>type</code>. <code>arg</code> is currently not used, although a future type may define new behavior parameters by way of <code>arg</code>.</p> <p><code>type</code> may be one of the following:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>USYNC_PROCESS</code></td> <td>The readers/writer lock can synchronize threads in this process and other processes. The readers/writer lock should be initialized by only one process. <code>arg</code> is ignored. A readers/writer lock initialized with this type, must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>USYNC_THREAD</code></td> <td>The readers/writer lock can synchronize threads in this process, only. <code>arg</code> is ignored.</td> </tr> </table> <p>Additionally, readers/writer locks can be initialized by allocation in zeroed memory. A type of <code>USYNC_THREAD</code> is assumed in this case. Multiple threads must not</p>	<code>USYNC_PROCESS</code>	The readers/writer lock can synchronize threads in this process and other processes. The readers/writer lock should be initialized by only one process. <code>arg</code> is ignored. A readers/writer lock initialized with this type, must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory.	<code>USYNC_THREAD</code>	The readers/writer lock can synchronize threads in this process, only. <code>arg</code> is ignored.
<code>USYNC_PROCESS</code>	The readers/writer lock can synchronize threads in this process and other processes. The readers/writer lock should be initialized by only one process. <code>arg</code> is ignored. A readers/writer lock initialized with this type, must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory.				
<code>USYNC_THREAD</code>	The readers/writer lock can synchronize threads in this process, only. <code>arg</code> is ignored.				

rwlock(3THR)

simultaneously initialize the same readers/writer lock. And a readers/writer lock must not be re-initialized while in use by other threads.

The following are default readers/writer lock initialization (intra-process):

```
rwlock_t rwlp;  
rwlock_init(&rwlp, NULL, NULL);  
OR  
rwlock_init(&rwlp, USYNC_THREAD, NULL);  
OR  
rwlock_t rwlp = DEFAULTRWLOCK;
```

The following is a customized readers/writer lock initialization (inter-process):

```
rwlock_init(&rwlp, USYNC_PROCESS, NULL);
```

Any state associated with the readers/writer lock pointed to by *rwlp* are destroyed by `rwlock_destroy()` and the readers/writer lock storage space is not released.

`rw_rdlock()` gets a read lock on the readers/writer lock pointed to by *rwlp*. If the readers/writer lock is currently locked for writing, the calling thread blocks until the write lock is freed. Multiple threads may simultaneously hold a read lock on a readers/writer lock.

`rw_tryrdlock()` tries to get a read lock on the readers/writer lock pointed to by *rwlp*. If the readers/writer lock is locked for writing, it returns an error; otherwise, the read lock is acquired.

`rw_wrlock()` gets a write lock on the readers/writer lock pointed to by *rwlp*. If the readers/writer lock is currently locked for reading or writing, the calling thread blocks until all the read and write locks are freed. At any given time, only one thread may have a write lock on a readers/writer lock.

`rw_trywrlock()` tries to get a write lock on the readers/writer lock pointed to by *rwlp*. If the readers/writer lock is currently locked for reading or writing, it returns an error.

`rw_unlock()` unlocks a readers/writer lock pointed to by *rwlp*, if the readers/writer lock is locked and the calling thread holds the lock for either reading or writing. One of the other threads that is waiting for the readers/writer lock to be freed will be unblocked, provided there is other waiting threads. If the calling thread does not hold the lock for either reading or writing, no error status is returned, and the program's behavior is unknown.

RETURN VALUES If successful, these functions return 0. Otherwise, a non-zero value is returned to indicate the error.

ERRORS The `rwlock_init()` function will fail if:

`EINVAL` `type` is invalid.

The `rw_tryrdlock()` or `rw_trywrlock()` functions will fail if:

rwlock(3THR)

EBUSY The reader or writer lock pointed to by *rwlp* was already locked.

These functions may fail if:

EFAULT *rwlp* or *arg* points to an illegal address.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `mmap(2)`, `attributes(5)`

NOTES These interfaces also available by way of:

```
#include <thread.h>
```

If multiple threads are waiting for a readers/writer lock, the acquisition order is random by default. However, some implementations may bias acquisition order to avoid depriving writers. The current implementation favors writers over readers.

schedctl_init(3SCHED)

NAME	<code>schedctl_init</code> , <code>schedctl_lookup</code> , <code>schedctl_exit</code> , <code>schedctl_start</code> , <code>schedctl_stop</code> – preemption control
SYNOPSIS	<pre>cc [flag ...] file ... -lsched [library ...] #include <schedctl.h> schedctl_t *schedctl_init(void); schedctl_t *schedctl_lookup(void); void schedctl_exit(void); void schedctl_start(schedctl_t *ptr); void schedctl_stop(schedctl_t *ptr);</pre>
DESCRIPTION	<p>These functions provide limited control over the scheduling of a lightweight process (LWP). They allow a running LWP to give a hint to the kernel that preemptions of that LWP should be avoided. The most likely use for these functions is to block preemption while holding a spinlock. Improper use of this facility, including attempts to block preemption for sustained periods of time, may result in reduced performance.</p> <p>The <code>schedctl_init()</code> function initializes preemption control for the calling LWP and returns a pointer used to refer to the data. If <code>schedctl_init()</code> is called more than once by the same LWP, the most recently returned pointer is the only valid one.</p> <p>The <code>schedctl_lookup()</code> function returns the currently allocated preemption control data associated with the calling LWP that was previously returned by <code>schedctl_init()</code>. This can be useful in programs where it is difficult to maintain local state for each LWP.</p> <p>The <code>schedctl_exit()</code> function removes the preemption control data associated with the calling LWP.</p> <p>The <code>schedctl_start()</code> macro gives a hint to the kernel scheduler that preemption should be avoided on the current LWP. The pointer passed to the macro must be the same as the pointer returned by the call to <code>schedctl_init()</code> by the current LWP. The behavior of the program when other values are passed is undefined.</p> <p>The <code>schedctl_stop()</code> macro removes the hint that was set by <code>schedctl_start()</code>. As with <code>schedctl_start()</code>, the pointer passed to the macro must be the same as the pointer returned by the call to <code>schedctl_init()</code> by the current LWP.</p> <p>The <code>schedctl_start()</code> and <code>schedctl_stop()</code> macros are intended to be used to bracket short critical sections, such as the time spent holding a spinlock. Other uses, including the failure to call <code>schedctl_stop()</code> soon after calling <code>schedctl_start()</code>, might result in poor performance.</p>

schedctl_init(3SCHED)

RETURN VALUES	The <code>schedctl_init()</code> function returns a pointer to a <code>schedctl_t</code> structure if the initialization was successful, or <code>NULL</code> otherwise. The <code>schedctl_lookup()</code> function returns a pointer to a <code>schedctl_t</code> structure if the data for that LWP was found, or <code>NULL</code> otherwise.
ERRORS	No errors are returned.
SEE ALSO	<code>priocntl(1)</code> , <code>exec(2)</code> , <code>fork(2)</code> , <code>priocntl(2)</code> , <code>thr_create(3THR)</code>
NOTES	<p>Preemption control is intended for use by LWPs belonging to the time-sharing (TS), interactive (IA), fair-share (FSS), and fixed-priority (FX) scheduling classes. If used by LWPs in other scheduling classes, such as real-time (RT), no errors will be returned but <code>schedctl_start()</code> and <code>schedctl_stop()</code> will not have any effect.</p> <p>Use of preemption control by unbound threads in multithreaded applications (see <code>thr_create(3THR)</code>) is not supported and will result in undefined behavior.</p> <p>The data used for preemption control are not copied in the child of a <code>fork(2)</code>. Thus, if a process containing LWPs using preemption control calls <code>fork</code> and the child does not immediately call <code>exec(2)</code>, each LWP in the child must call <code>schedctl_init()</code> again prior to any future uses of <code>schedctl_start()</code> and <code>schedctl_stop()</code>. Failure to do so will result in undefined behavior.</p>

sched_getparam(3RT)

NAME	sched_getparam – get scheduling parameters						
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <sched.h> int sched_getparam(pid_t pid, struct sched_param *param);</pre>						
DESCRIPTION	<p>The sched_getparam() function returns the scheduling parameters of a process specified by <i>pid</i> in the sched_param structure pointed to by <i>param</i>.</p> <p>If a process specified by <i>pid</i> exists and if the calling process has permission, the scheduling parameters for the process whose process ID is equal to <i>pid</i> will be returned.</p> <p>If <i>pid</i> is 0, the scheduling parameters for the calling process will be returned. The behavior of the sched_getparam() function is unspecified if the value of <i>pid</i> is negative.</p>						
RETURN VALUES	Upon successful completion, the sched_getparam() function returns 0. If the call to sched_getparam() is unsuccessful, the function returns -1 and sets errno to indicate the error.						
ERRORS	The sched_getparam() function will fail if: ENOSYS The sched_getparam() function is not supported by the system. EPERM The requesting process does not have permission to obtain the scheduling parameters of the specified process. ESRCH No process can be found corresponding to that specified by <i>pid</i> .						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	librt(3LIB), sched(3HEAD), sched_getscheduler(3RT), sched_setparam(3RT), sched_setscheduler(3RT), attributes(5)						
NOTES	Solaris 2.6 was the first release to support libposix4/librt. Prior to this release, this function always returned -1 and set errno to ENOSYS.						

sched_get_priority_max(3RT)

NAME sched_get_priority_max, sched_get_priority_min – get scheduling parameter limits

SYNOPSIS

```
cc [ flag... ] file... -lrt [ library... ]
#include <sched.h>

int sched_get_priority_max(int policy);
int sched_get_priority_min(int policy);
```

DESCRIPTION

The sched_get_priority_max() and sched_get_priority_min() functions return the appropriate maximum or minimum, respectfully, for the scheduling policy specified by *policy*.

The value of *policy* is one of the scheduling policy values defined in <sched.h>.

RETURN VALUES

If successful, the sched_get_priority_max() and sched_get_priority_min() functions return the appropriate maximum or minimum values, respectively. If unsuccessful, they return -1 and set errno to indicate the error.

ERRORS

The sched_get_priority_max() and sched_get_priority_min() functions will fail if:

EINVAL	The value of the <i>policy</i> parameter does not represent a defined scheduling policy.
ENOSYS	The sched_get_priority_max(), sched_get_priority_min() and sched_rr_get_interval(3RT) functions are not supported by the system.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO librt(3LIB), sched(3HEAD), sched_getparam(3RT), sched_setparam(3RT), sched_getscheduler(3RT), sched_rr_get_interval(3RT), sched_setscheduler(3RT), time(3HEAD), attributes(5)

NOTES

Solaris 2.6 was the first release to support libposix4/librt. Prior to this release, this function always returned -1 and set errno to ENOSYS.

sched_getscheduler(3RT)

NAME	sched_getscheduler – get scheduling policy						
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <sched.h> int sched_getscheduler(pid_t pid);</pre>						
DESCRIPTION	<p>The <code>sched_getscheduler()</code> function returns the scheduling policy of the process specified by <i>pid</i>. If the value of <i>pid</i> is negative, the behavior of the <code>sched_getscheduler()</code> function is unspecified.</p> <p>The values that can be returned by <code>sched_getscheduler()</code> are defined in the header <code><sched.h></code> and described on the <code>sched_setscheduler(3RT)</code> manual page.</p> <p>If a process specified by <i>pid</i> exists and if the calling process has permission, the scheduling policy will be returned for the process whose process ID is equal to <i>pid</i>.</p> <p>If <i>pid</i> is 0, the scheduling policy will be returned for the calling process.</p>						
RETURN VALUES	Upon successful completion, the <code>sched_getscheduler()</code> function returns the scheduling policy of the specified process. If unsuccessful, the function returns <code>-1</code> and sets <code>errno</code> to indicate the error.						
ERRORS	The <code>sched_getscheduler()</code> function will fail if: ENOSYS The <code>sched_getscheduler()</code> function is not supported by the system. EPERM The requesting process does not have permission to determine the scheduling policy of the specified process. ESRCH No process can be found corresponding to that specified by <i>pid</i> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>librt(3LIB)</code> , <code>sched(3HEAD)</code> , <code>sched_getparam(3RT)</code> , <code>sched_setparam(3RT)</code> , <code>sched_setscheduler(3RT)</code> , <code>attributes(5)</code>						
NOTES	Solaris 2.6 was the first release to support <code>libposix4/librt</code> . Prior to this release, this function always returned <code>-1</code> and set <code>errno</code> to <code>ENOSYS</code> .						

NAME sched_rr_get_interval – get execution time limits

SYNOPSIS

```
cc [ flag... ] file... -lrt [ library... ]
#include <sched.h>

int sched_rr_get_interval(pid_t pid, struct timespec *interval);
```

DESCRIPTION The sched_rr_get_interval() function updates the timespec structure referenced by the *interval* argument to contain the current execution time limit (that is, time quantum) for the process specified by *pid*. If *pid* is 0, the current execution time limit for the calling process will be returned.

RETURN VALUES If successful, the sched_rr_get_interval() function returns 0. Otherwise, it returns -1 and sets errno to indicate the error.

ERRORS The sched_rr_get_interval() function will fail if:

- ENOSYS The sched_get_priority_max(3RT), sched_get_priority_min(3RT), and sched_rr_get_interval() functions are not supported by the system.
- ESRCH No process can be found corresponding to that specified by *pid*.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO librt(3LIB), sched(3HEAD), sched_getparam(3RT), sched_setparam(3RT), sched_get_priority_max(3RT), sched_getscheduler(3RT), sched_setscheduler(3RT), attributes(5)

NOTES Solaris 2.6 was the first release to support libposix4/librt. Prior to this release, this function always returned -1 and set errno to ENOSYS.

sched_setparam(3RT)

NAME	sched_setparam – set scheduling parameters
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <sched.h> int sched_setparam(pid_t pid, const struct sched_param *param);</pre>
DESCRIPTION	<p>The <code>sched_setparam()</code> function sets the scheduling parameters of the process specified by <code>pid</code> to the values specified by the <code>sched_param</code> structure pointed to by <code>param</code>. The value of the <code>sched_priority</code> member in the <code>sched_param</code> structure is any integer within the inclusive priority range for the current scheduling policy of the process specified by <code>pid</code>. Higher numerical values for the priority represent higher priorities. If the value of <code>pid</code> is negative, the behavior of the <code>sched_setparam()</code> function is unspecified.</p> <p>If a process specified by <code>pid</code> exists and if the calling process has permission, the scheduling parameters will be set for the process whose process ID is equal to <code>pid</code>. The real or effective user ID of the calling process must match the real or saved (from <code>exec(2)</code>) user ID of the target process unless the effective user ID of the calling process is 0. See <code>intro(2)</code>.</p> <p>If <code>pid</code> is zero, the scheduling parameters will be set for the calling process.</p> <p>The target process, whether it is running or not running, resumes execution after all other runnable processes of equal or greater priority have been scheduled to run.</p> <p>If the priority of the process specified by the <code>pid</code> argument is set higher than that of the lowest priority running process and if the specified process is ready to run, the process specified by the <code>pid</code> argument preempts a lowest priority running process. Similarly, if the process calling <code>sched_setparam()</code> sets its own priority lower than that of one or more other non-empty process lists, then the process that is the head of the highest priority list also preempts the calling process. Thus, in either case, the originating process might not receive notification of the completion of the requested priority change until the higher priority process has executed.</p> <p>If the current scheduling policy for the process specified by <code>pid</code> is not <code>SCHED_FIFO</code> or <code>SCHED_RR</code>, including <code>SCHED_OTHER</code>, the result is equal to <code>prctl(P_PID, pid, PC_SETPARMS, &pcparam)</code>, where <code>pcparam</code> is an image of <code>*param</code>.</p> <p>The effect of this function on individual threads is dependent on the scheduling contention scope of the threads:</p> <ul style="list-style-type: none">■ For threads with system scheduling contention scope, these functions have no effect on their scheduling.■ For threads with process scheduling contention scope, the threads' scheduling parameters will not be affected. However, the scheduling of these threads with respect to threads in other processes may be dependent on the scheduling parameters of their process, which are governed using these functions.

If an implementation supports a two-level scheduling model in which library threads are multiplexed on top of several kernel scheduled entities, then the underlying kernel scheduled entities for the system contention scope threads will not be affected by these functions.

The underlying kernel scheduled entities for the process contention scope threads will have their scheduling parameters changed to the value specified in *param*. Kernel scheduled entities for use by process contention scope threads that are created after this call completes inherit their scheduling policy and associated scheduling parameters from the process.

This function is not atomic with respect to other threads in the process. Threads are allowed to continue to execute while this function call is in the process of changing the scheduling policy for the underlying kernel scheduled entities used by the process contention scope threads.

RETURN VALUES

If successful, the `sched_setparam()` function returns 0.

If the call to `sched_setparam()` is unsuccessful, the priority remains unchanged, and the function returns -1 and sets `errno` to indicate the error.

ERRORS

The `sched_setparam()` function will fail if:

- EINVAL One or more of the requested scheduling parameters is outside the range defined for the scheduling policy of the specified *pid*.
- ENOSYS The `sched_setparam()` function is not supported by the system.
- EPERM The requesting process does not have permission to set the scheduling parameters for the specified process, or does not have the appropriate privilege to invoke `sched_setparam()`.
- ESRCH No process can be found corresponding to that specified by *pid*.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

`intro(2)`, `exec(2)`, `librt(3LIB)`, `sched(3HEAD)`, `sched_getparam(3RT)`, `sched_getscheduler(3RT)`, `sched_setscheduler(3RT)`, `attributes(5)`

NOTES

Solaris 2.6 was the first release to support `libposix4/librt`. Prior to this release, this function always returned -1 and set `errno` to `ENOSYS`.

sched_setscheduler(3RT)

NAME	sched_setscheduler – set scheduling policy and scheduling parameters
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <sched.h> int sched_setscheduler(pid_t pid, int policy, const struct sched_param *param);</pre>
DESCRIPTION	<p>The <code>sched_setscheduler()</code> function sets the scheduling policy and scheduling parameters of the process specified by <i>pid</i> to <i>policy</i> and the parameters specified in the <code>sched_param</code> structure pointed to by <i>param</i>, respectively. The value of the <code>sched_priority</code> member in the <code>sched_param</code> structure is any integer within the inclusive priority range for the scheduling policy specified by <i>policy</i>. The <code>sched_setscheduler()</code> function ignores the other members of the <code>sched_param</code> structure. If the value of <i>pid</i> is negative, the behavior of the <code>sched_setscheduler()</code> function is unspecified.</p> <p>The possible values for the <i>policy</i> parameter are defined in the header <code><sched.h></code> (see <code>sched(3HEAD)</code>):</p> <p>If a process specified by <i>pid</i> exists and if the calling process has permission, the scheduling policy and scheduling parameters are set for the process whose process ID is equal to <i>pid</i>. The real or effective user ID of the calling process must match the real or saved (from <code>exec(2)</code>) user ID of the target process unless the effective user ID of the calling process is 0. See <code>intro(2)</code>.</p> <p>If <i>pid</i> is 0, the scheduling policy and scheduling parameters are set for the calling process.</p> <p>To change the <i>policy</i> of any process to either of the real time policies <code>SCHED_FIFO</code> or <code>SCHED_RR</code>, the calling process must either have the <code>SCHED_FIFO</code> or <code>SCHED_RR</code> policy or have an effective user ID of 0.</p> <p>The <code>sched_setscheduler()</code> function is considered successful if it succeeds in setting the scheduling policy and scheduling parameters of the process specified by <i>pid</i> to the values specified by <i>policy</i> and the structure pointed to by <i>param</i>, respectively.</p> <p>The effect of this function on individual threads is dependent on the scheduling contention scope of the threads:</p> <ul style="list-style-type: none">■ For threads with system scheduling contention scope, these functions have no effect on their scheduling.■ For threads with process scheduling contention scope, the threads' scheduling policy and associated parameters will not be affected. However, the scheduling of these threads with respect to threads in other processes may be dependent on the scheduling parameters of their process, which are governed using these functions.

The system supports a two-level scheduling model in which library threads are multiplexed on top of several kernel scheduled entities. The underlying kernel scheduled entities for the system contention scope threads will not be affected by these functions.

The underlying kernel scheduled entities for the process contention scope threads will have their scheduling policy and associated scheduling parameters changed to the values specified in *policy* and *param*, respectively. Kernel scheduled entities for use by process contention scope threads that are created after this call completes inherit their scheduling policy and associated scheduling parameters from the process.

This function is not atomic with respect to other threads in the process. Threads are allowed to continue to execute while this function call is in the process of changing the scheduling policy and associated scheduling parameters for the underlying kernel scheduled entities used by the process contention scope threads.

RETURN VALUES

Upon successful completion, the function returns the former scheduling policy of the specified process. If the `sched_setscheduler()` function fails to complete successfully, the policy and scheduling parameters remain unchanged, and the function returns `-1` and sets `errno` to indicate the error.

ERRORS

The `sched_setscheduler()` function will fail if:

- EINVAL** The value of *policy* is invalid, or one or more of the parameters contained in *param* is outside the valid range for the specified scheduling policy.
- ENOSYS** The `sched_setscheduler()` function is not supported by the system.
- EPERM** The requesting process does not have permission to set either or both of the scheduling parameters or the scheduling policy of the specified process.
- ESRCH** No process can be found corresponding to that specified by *pid*.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

`priocntl(1)`, `intro(2)`, `exec(2)`, `priocntl(2)`, `librt(3LIB)`, `sched(3HEAD)`, `sched_get_priority_max(3RT)`, `sched_getparam(3RT)`, `sched_getscheduler(3RT)`, `sched_setparam(3RT)`, `attributes(5)`

NOTES

Solaris 2.6 was the first release to support `libposix4/librt`. Prior to this release, this function always returned `-1` and set `errno` to `ENOSYS`.

sched_yield(3RT)

NAME sched_yield – yield processor

SYNOPSIS

```
cc [ flag... ] file... -lrt [ library... ]
#include <sched.h>

int sched_yield(void);
```

DESCRIPTION The sched_yield() function forces the running thread to relinquish the processor until the process again becomes the head of its process list. It takes no arguments.

RETURN VALUES If successful, sched_yield() returns 0, otherwise, it returns -1, and sets errno to indicate the error condition.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO librt(3LIB), sched(3HEAD), attributes(5)

NAME	semaphore, sema_init, sema_destroy, sema_wait, sema_trywait, sema_post – semaphores				
SYNOPSIS	<pre>cc [<i>flag...</i>] <i>file...</i>- lthread - lc [<i>library...</i>] #include <synch.h> int sema_init(sema_t *sp, unsigned int <i>count</i>, int <i>type</i>, void * <i>arg</i>); int sema_destroy(sema_t *sp); int sema_wait(sema_t *sp); int sema_trywait(sema_t *sp); int sema_post(sema_t *sp);</pre>				
DESCRIPTION	<p>A semaphore is a non-negative integer count and is generally used to coordinate access to resources. The initial semaphore count is set to the number of free resources, then threads slowly increment and decrement the count as resources are added and removed. If the semaphore count drops to zero, which means no available resources, threads attempting to decrement the semaphore will block until the count is greater than zero.</p> <p>Semaphores can synchronize threads in this process and other processes if they are allocated in writable memory and shared among the cooperating processes (see <code>mmap(2)</code>), and have been initialized for this purpose.</p> <p>Semaphores must be initialized before use; semaphores pointed to by <i>sp</i> to <i>count</i> are initialized by <code>sema_init()</code>. The <i>type</i> argument can assign several different types of behavior to a semaphore. No current type uses <i>arg</i>, although it may be used in the future.</p> <p>The <i>type</i> argument may be one of the following:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">USYNC_PROCESS</td> <td>The semaphore can synchronize threads in this process and other processes. Initializing the semaphore should be done by only one process. A semaphore initialized with this type must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>), or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory. <i>arg</i> is ignored.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">USYNC_THREAD</td> <td>The semaphore can synchronize threads only in this process. The <i>arg</i> argument is ignored. USYNC_THREAD does not support multiple mappings to the same logical synch object. If you need to <code>mmap()</code> a synch object to different locations within the same address space, then the synch object should be initialized as a shared object USYNC_PROCESS for Solaris threads and</td> </tr> </table>	USYNC_PROCESS	The semaphore can synchronize threads in this process and other processes. Initializing the semaphore should be done by only one process. A semaphore initialized with this type must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>), or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory. <i>arg</i> is ignored.	USYNC_THREAD	The semaphore can synchronize threads only in this process. The <i>arg</i> argument is ignored. USYNC_THREAD does not support multiple mappings to the same logical synch object. If you need to <code>mmap()</code> a synch object to different locations within the same address space, then the synch object should be initialized as a shared object USYNC_PROCESS for Solaris threads and
USYNC_PROCESS	The semaphore can synchronize threads in this process and other processes. Initializing the semaphore should be done by only one process. A semaphore initialized with this type must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>), or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory. <i>arg</i> is ignored.				
USYNC_THREAD	The semaphore can synchronize threads only in this process. The <i>arg</i> argument is ignored. USYNC_THREAD does not support multiple mappings to the same logical synch object. If you need to <code>mmap()</code> a synch object to different locations within the same address space, then the synch object should be initialized as a shared object USYNC_PROCESS for Solaris threads and				

semaphore(3THR)

PTHREAD_PROCESS_PRIVATE for POSIX threads.

A semaphore must not be simultaneously initialized by multiple threads, nor re-initialized while in use by other threads.

Default semaphore initialization (intra-process):

```
sema_t sp;  
int count = 1;  
sema_init(&sp, count, NULL, NULL);
```

or

```
sema_init(&sp, count, USYNC_THREAD, NULL);
```

Customized semaphore initialization (inter-process):

```
sema_t sp;  
int count = 1;  
sema_init(&sp, count, USYNC_PROCESS, NULL);
```

The `sema_destroy()` function destroys any state related to the semaphore pointed to by `sp`. The semaphore storage space is not released.

The `sema_wait()` function blocks the calling thread until the semaphore count pointed to by `sp` is greater than zero, and then it atomically decrements the count.

The `sema_trywait()` function atomically decrements the semaphore count pointed to by `sp`, if the count is greater than zero; otherwise, it returns an error.

The `sema_post()` function atomically increments the semaphore count pointed to by `sp`. If there are any threads blocked on the semaphore, one will be unblocked.

The semaphore functionality described on this man page is for the Solaris threads implementation. For the POSIX-compliant semaphore interface documentation, see `sem_open(3RT)`, `sem_init(3RT)`, `sem_wait(3RT)`, `sem_post(3RT)`, `sem_getvalue(3RT)`, `sem_unlink(3RT)`, `sem_close(3RT)`, `sem_destroy(3RT)`.

RETURN VALUES

Upon successful completion, 0 is returned; otherwise, a non-zero value indicates an error.

ERRORS

These functions will fail if:

EINVAL	The <code>sp</code> argument does not refer to a valid semaphore..
EFAULT	Either the <code>sp</code> or <code>arg</code> argument points to an illegal address.

The `sema_wait()` function will fail if::

EINTR	The wait was interrupted by a signal or <code>fork()</code> .
-------	---

The `sema_trywait()` function will fail if::

EBUSY	The semaphore pointed to by <code>sp</code> has a zero count.
-------	---

The `sema_post()` function will fail if:

`EOVERFLOW` The semaphore value pointed to by `sp` exceeds `SEM_VALUE_MAX`.

EXAMPLES

EXAMPLE 1 The customer waiting-line in a bank is analogous to the synchronization scheme of a semaphore using `sema_wait()` and `sema_trywait()`:

```
/* cc [ flag . . . ] file . . . -lthread [ library . . . ] */
#include <errno.h>
#define TELLERS 10
sema_t    tellers;    /* semaphore */
int banking_hours(), deposit_withdrawal;
void*customer(), do_business(), skip_banking_today();
. . .

sema_init(&tellers, TELLERS, USYNC_THREAD, NULL);
/* 10 tellers available */
while(banking_hours())
    pthread_create(NULL, NULL, customer, deposit_withdrawal);
. . .

void *
customer(int deposit_withdrawal)
{
    int this_customer, in_a_hurry = 50;
    this_customer = rand() % 100;

    if (this_customer == in_a_hurry) {
        if (sema_trywait(&tellers) != 0)
            if (errno == EAGAIN) { /* no teller available */
                skip_banking_today(this_customer);
                return;
            } /* else go immediately to available teller and
                decrement tellers */
    }
    else
        sema_wait(&tellers); /* wait for next teller, then proceed,
                            and decrement tellers */

    do_business(deposit_withdrawal);
    sema_post(&tellers); /* increment tellers;
                        this_customer's teller
                        is now available */
}
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

semaphore(3THR)

SEE ALSO | mmap(2), shmop(2), sem_close(3RT), sem_destroy(3RT), sem_getvalue(3RT),
sem_init(3RT), sem_open(3RT), sem_post(3RT), sem_unlink(3RT),
sem_wait(3RT), attributes(5), standards(5)

NOTES | These functions are also available by way of:

```
#include <thread.h>
```

By default, there is no defined order of unblocking for multiple threads waiting for a semaphore.

NAME	sem_close – close a named semaphore				
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <semaphore.h> int sem_close(sem_t *sem);</pre>				
DESCRIPTION	<p>The <code>sem_close()</code> function is used to indicate that the calling process is finished using the named semaphore indicated by <code>sem</code>. The effects of calling <code>sem_close()</code> for an unnamed semaphore (one created by <code>sem_init(3RT)</code>) are undefined. The <code>sem_close()</code> function deallocates (that is, make available for reuse by a subsequent <code>sem_open(3RT)</code> by this process) any system resources allocated by the system for use by this process for this semaphore. The effect of subsequent use of the semaphore indicated by <code>sem</code> by this process is undefined. If the semaphore has not been removed with a successful call to <code>sem_unlink(3RT)</code>, then <code>sem_close()</code> has no effect on the state of the semaphore. If the <code>sem_unlink(3RT)</code> function has been successfully invoked for <code>name</code> after the most recent call to <code>sem_open(3RT)</code> with <code>O_CREAT</code> for this semaphore, then when all processes that have opened the semaphore close it, the semaphore is no longer be accessible.</p>				
RETURN VALUES	If successful, <code>sem_close()</code> returns 0, otherwise it returns -1 and sets <code>errno</code> to indicate the error.				
ERRORS	<p>The <code>sem_close()</code> function will fail if:</p> <p><code>EINVAL</code> The <code>sem</code> argument is not a valid semaphore descriptor.</p> <p><code>ENOSYS</code> The <code>sem_close()</code> function is not supported by the system.</p>				
USAGE	The <code>sem_close()</code> function should not be called for an unnamed semaphore initialized by <code>sem_init(3RT)</code> .				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>sem_init(3RT)</code> , <code>sem_open(3RT)</code> , <code>sem_unlink(3RT)</code> , <code>attributes(5)</code>				
NOTES	Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set <code>errno</code> to <code>ENOSYS</code> .				

sem_destroy(3RT)

NAME sem_destroy – destroy an unnamed semaphore

SYNOPSIS

```
cc [ flag... ] file... -lrt [ library... ]
#include <semaphore.h>

int sem_destroy(sem_t *sem);
```

DESCRIPTION The `sem_destroy()` function is used to destroy the unnamed semaphore indicated by `sem`. Only a semaphore that was created using `sem_init(3RT)` may be destroyed using `sem_destroy()`; the effect of calling `sem_destroy()` with a named semaphore is undefined. The effect of subsequent use of the semaphore `sem` is undefined until `sem` is re-initialized by another call to `sem_init(3RT)`.

It is safe to destroy an initialised semaphore upon which no threads are currently blocked. The effect of destroying a semaphore upon which other threads are currently blocked is undefined.

RETURN VALUES If successful, `sem_destroy()` returns 0, otherwise it returns -1 and sets `errno` to indicate the error.

ERRORS The `sem_destroy()` function will fail if:

- EINVAL The `sem` argument is not a valid semaphore.
- ENOSYS The `sem_destroy()` function is not supported by the system.

The `sem_destroy()` function may fail if:

- EBUSY There are currently processes (or LWPs or threads) blocked on the semaphore.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `sem_init(3RT)`, `sem_open(3RT)`, `attributes(5)`

NAME	sem_getvalue – get the value of a semaphore				
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <semaphore.h> int sem_getvalue(sem_t *sem, int *sval);</pre>				
DESCRIPTION	<p>The <code>sem_getvalue()</code> function updates the location referenced by the <code>sval</code> argument to have the value of the semaphore referenced by <code>sem</code> without affecting the state of the semaphore. The updated value represents an actual semaphore value that occurred at some unspecified time during the call, but it need not be the actual value of the semaphore when it is returned to the calling process.</p> <p>If <code>sem</code> is locked, then the value returned by <code>sem_getvalue()</code> is either zero or a negative number whose absolute value represents the number of processes waiting for the semaphore at some unspecified time during the call.</p> <p>The value set in <code>sval</code> may be 0 or positive. If <code>sval</code> is 0, there may be other processes (or LWPs or threads) waiting for the semaphore; if <code>sval</code> is positive, no processed is waiting.</p>				
RETURN VALUES	Upon successful completion, <code>sem_getvalue()</code> returns 0. Otherwise, it returns <code>-1</code> and sets <code>errno</code> to indicate the error.				
ERRORS	<p>The <code>sem_getvalue()</code> function will fail if:</p> <p><code>EINVAL</code> The <code>sem</code> argument does not refer to a valid semaphore.</p> <p><code>ENOSYS</code> The <code>sem_getvalue()</code> function is not supported by the system.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>sem_post(3RT)</code> , <code>sem_wait(3RT)</code> , <code>attributes(5)</code>				

sem_init(3RT)

NAME	sem_init – initialize an unnamed semaphore
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <semaphore.h> int sem_init(sem_t *sem, int pshared, unsigned int value);</pre>
DESCRIPTION	<p>The <code>sem_init()</code> function is used to initialize the unnamed semaphore referred to by <code>sem</code>. The value of the initialized semaphore is <code>value</code>. Following a successful call to <code>sem_init()</code>, the semaphore may be used in subsequent calls to <code>sem_wait(3RT)</code>, <code>sem_trywait(3RT)</code>, <code>sem_post(3RT)</code>, and <code>sem_destroy(3RT)</code>. This semaphore remains usable until the semaphore is destroyed.</p> <p>If the <code>pshared</code> argument has a non-zero value, then the semaphore is shared between processes; in this case, any process that can access the semaphore <code>sem</code> can use <code>sem</code> for performing <code>sem_wait(3RT)</code>, <code>sem_trywait(3RT)</code>, <code>sem_post(3RT)</code>, and <code>sem_destroy(3RT)</code> operations.</p> <p>Only <code>sem</code> itself may be used for performing synchronization. The result of referring to copies of <code>sem</code> in calls to <code>sem_wait(3RT)</code>, <code>sem_trywait(3RT)</code>, <code>sem_post(3RT)</code>, and <code>sem_destroy(3RT)</code>, is undefined.</p> <p>If the <code>pshared</code> argument is zero, then the semaphore is shared between threads of the process; any thread in this process can use <code>sem</code> for performing <code>sem_wait(3RT)</code>, <code>sem_trywait(3RT)</code>, <code>sem_post(3RT)</code>, and <code>sem_destroy(3RT)</code> operations. The use of the semaphore by threads other than those created in the same process is undefined.</p> <p>Attempting to initialize an already initialized semaphore results in undefined behavior.</p>
RETURN VALUES	Upon successful completion, the function initializes the semaphore in <code>sem</code> . Otherwise, it returns <code>-1</code> and sets <code>errno</code> to indicate the error.
ERRORS	The <code>sem_init()</code> function will fail if:
EINVAL	The <code>value</code> argument exceeds <code>SEM_VALUE_MAX</code> .
ENOSPC	A resource required to initialize the semaphore has been exhausted, or the resources have reached the limit on semaphores (<code>SEM_NSEMS_MAX</code>).
ENOSYS	The <code>sem_init()</code> function is not supported by the system.
EPERM	The process lacks the appropriate privileges to initialize the semaphore.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

`sem_init(3RT)`

SEE ALSO `sem_destroy(3RT)`, `sem_post(3RT)`, `sem_wait(3RT)`, `attributes(5)`

sem_open(3RT)

NAME	sem_open – initialize/open a named semaphore
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <semaphore.h> sem_t *sem_open(const char *name, int oflag, /* unsigned long mode, unsigned int value */ ...);</pre>
DESCRIPTION	<p>The <code>sem_open()</code> function establishes a connection between a named semaphore and a process (or LWP or thread). Following a call to <code>sem_open()</code> with semaphore name <i>name</i>, the process may reference the semaphore associated with <i>name</i> using the address returned from the call. This semaphore may be used in subsequent calls to <code>sem_wait(3RT)</code>, <code>sem_trywait(3RT)</code>, <code>sem_post(3RT)</code>, and <code>sem_close(3RT)</code>. The semaphore remains usable by this process until the semaphore is closed by a successful call to <code>sem_close(3RT)</code>, <code>_exit(2)</code>, or one of the <code>exec</code> functions.</p> <p>The <i>oflag</i> argument controls whether the semaphore is created or merely accessed by the call to <code>sem_open()</code>. The following flag bits may be set in <i>oflag</i>:</p> <p>O_CREAT This flag is used to create a semaphore if it does not already exist. If <code>O_CREAT</code> is set and the semaphore already exists, then <code>O_CREAT</code> has no effect, except as noted under <code>O_EXCL</code>. Otherwise, <code>sem_open()</code> creates a named semaphore. The <code>O_CREAT</code> flag requires a third and a fourth argument: <i>mode</i>, which is of type <code>mode_t</code>, and <i>value</i>, which is of type <code>unsigned int</code>. The semaphore is created with an initial value of <i>value</i>. Valid initial values for semaphores are less than or equal to <code>SEM_VALUE_MAX</code>.</p> <p>The user ID of the semaphore is set to the effective user ID of the process; the group ID of the semaphore is set to a system default group ID or to the effective group ID of the process. The permission bits of the semaphore are set to the value of the <i>mode</i> argument except those set in the file mode creation mask of the process (see <code>umask(2)</code>). When bits in <i>mode</i> other than the file permission bits are specified, the effect is unspecified.</p> <p>After the semaphore named <i>name</i> has been created by <code>sem_open()</code> with the <code>O_CREAT</code> flag, other processes can connect to the semaphore by calling <code>sem_open()</code> with the same value of <i>name</i>.</p> <p>O_EXCL If <code>O_EXCL</code> and <code>O_CREAT</code> are set, <code>sem_open()</code> fails if the semaphore <i>name</i> exists. The check for the existence of the semaphore and the creation of the semaphore if it does not exist are atomic with respect to other processes executing <code>sem_open()</code> with <code>O_EXCL</code> and <code>O_CREAT</code> set. If <code>O_EXCL</code> is set and <code>O_CREAT</code> is not set, the effect is undefined.</p> <p>If flags other than <code>O_CREAT</code> and <code>O_EXCL</code> are specified in the <i>oflag</i> parameter, the effect is unspecified.</p>

The *name* argument points to a string naming a semaphore object. It is unspecified whether the name appears in the file system and is visible to functions that take pathnames as arguments. The *name* argument conforms to the construction rules for a pathname. The first character of *name* must be a slash (/) character and the remaining characters of *name* cannot include any slash characters. For maximum portability, *name* should include no more than 14 characters, but this limit is not enforced.

If a process makes multiple successful calls to `sem_open()` with the same value for *name*, the same semaphore address is returned for each such successful call, provided that there have been no calls to `sem_unlink(3RT)` for this semaphore.

References to copies of the semaphore produce undefined results.

RETURN VALUES

Upon successful completion, the function returns the address of the semaphore. Otherwise, it will return a value of `SEM_FAILED` and set `errno` to indicate the error. The symbol `SEM_FAILED` is defined in the header `<semaphore.h>`. No successful return from `sem_open()` will return the value `SEM_FAILED`.

ERRORS

If any of the following conditions occur, the `sem_open()` function will return `SEM_FAILED` and set `errno` to the corresponding value:

<code>EACCES</code>	The named semaphore exists and the <code>O_RDWR</code> permissions are denied, or the named semaphore does not exist and permission to create the named semaphore is denied.
<code>EEXIST</code>	<code>O_CREAT</code> and <code>O_EXCL</code> are set and the named semaphore already exists.
<code>EINTR</code>	The <code>sem_open()</code> function was interrupted by a signal.
<code>EINVAL</code>	The <code>sem_open()</code> operation is not supported for the given name, or <code>O_CREAT</code> was set in <i>oflag</i> and <i>value</i> is greater than <code>SEM_VALUE_MAX</code> .
<code>EMFILE</code>	The number of open semaphore descriptors in this process exceeds <code>SEM_NSEMS_MAX</code> , or the number of open file descriptors in this process exceeds <code>OPEN_MAX</code> .
<code>ENAMETOOLONG</code>	The length of <i>name</i> string exceeds <code>PATH_MAX</code> , or a pathname component is longer than <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
<code>ENFILE</code>	Too many semaphores are currently open in the system.
<code>ENOENT</code>	<code>O_CREAT</code> is not set and the named semaphore does not exist.
<code>ENOSPC</code>	There is insufficient space for the creation of the new named semaphore.

sem_open(3RT)

ENOSYS

The `sem_open()` function is not supported by the system.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

`exec(2)`, `exit(2)`, `umask(2)`, `sem_close(3RT)`, `sem_post(3RT)`, `sem_unlink(3RT)`, `sem_wait(3RT)`, `sysconf(3C)`, `attributes(5)`

NAME	sem_post – increment the count of a semaphore
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <semaphore.h> int sem_post (sem_t *sem);</pre>
DESCRIPTION	<p>The sem_post () function unlocks the semaphore referenced by <i>sem</i> by performing a semaphore unlock operation on that semaphore.</p> <p>If the semaphore value resulting from this operation is positive, then no threads were blocked waiting for the semaphore to become unlocked; the semaphore value is simply incremented.</p> <p>If the value of the semaphore resulting from this operation is 0, then one of the threads blocked waiting for the semaphore will be allowed to return successfully from its call to sem_wait(3RT). If the symbol _POSIX_PRIORITY_SCHEDULING is defined, the thread to be unblocked will be chosen in a manner appropriate to the scheduling policies and parameters in effect for the blocked threads. In the case of the schedulers SCHED_FIFO and SCHED_RR, the highest priority waiting thread will be unblocked, and if there is more than one highest priority thread blocked waiting for the semaphore, then the highest priority thread that has been waiting the longest will be unblocked. If the symbol _POSIX_PRIORITY_SCHEDULING is not defined, the choice of a thread to unblock is unspecified.</p>
RETURN VALUES	If successful, sem_post () returns 0; otherwise it returns -1 and sets errno to indicate the error.
ERRORS	<p>The sem_post () function will fail if:</p> <p>EINVAL The <i>sem</i> argument does not refer to a valid semaphore.</p> <p>ENOSYS The sem_post () function is not supported by the system.</p> <p>EOVERFLOW The semaphore value exceeds SEM_VALUE_MAX.</p>
USAGE	The sem_post () function is reentrant with respect to signals and may be invoked from a signal-catching function. The semaphore functionality described on this manual page is for the POSIX (see standards(5)) threads implementation. For the documentation of the Solaris threads interface, see semaphore(3THR).
EXAMPLES	EXAMPLE 1 See sem_wait(3RT).
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

sem_post(3RT)

SEE ALSO | sched_setscheduler(3RT), sem_wait(3RT), semaphore(3THR), attributes(5), standards(5)

NAME	sem_unlink – remove a named semaphore								
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <semaphore.h> int sem_unlink(const char *name);</pre>								
DESCRIPTION	<p>The <code>sem_unlink()</code> function removes the semaphore named by the string <i>name</i>. If the semaphore named by <i>name</i> is currently referenced by other processes, then <code>sem_unlink()</code> has no effect on the state of the semaphore. If one or more processes have the semaphore open when <code>sem_unlink()</code> is called, destruction of the semaphore is postponed until all references to the semaphore have been destroyed by calls to <code>sem_close(3RT)</code>, <code>_exit(2)</code>, or one of the <code>exec</code> functions (see <code>exec(2)</code>). Calls to <code>sem_open(3RT)</code> to re-create or re-connect to the semaphore refer to a new semaphore after <code>sem_unlink()</code> is called. The <code>sem_unlink()</code> call does not block until all references have been destroyed; it returns immediately.</p>								
RETURN VALUES	<p>Upon successful completion, <code>sem_unlink()</code> returns 0. Otherwise, the semaphore is not changed and the function returns a value of -1 and sets <code>errno</code> to indicate the error.</p>								
ERRORS	<p>The <code>sem_unlink()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EACCES</td> <td>Permission is denied to unlink the named semaphore.</td> </tr> <tr> <td>ENAMETOOLONG</td> <td>The length of <i>name</i> string exceeds <code>PATH_MAX</code>, or a pathname component is longer than <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.</td> </tr> <tr> <td>ENOENT</td> <td>The named semaphore does not exist.</td> </tr> <tr> <td>ENOSYS</td> <td>The <code>sem_unlink()</code> function is not supported by the system.</td> </tr> </table>	EACCES	Permission is denied to unlink the named semaphore.	ENAMETOOLONG	The length of <i>name</i> string exceeds <code>PATH_MAX</code> , or a pathname component is longer than <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.	ENOENT	The named semaphore does not exist.	ENOSYS	The <code>sem_unlink()</code> function is not supported by the system.
EACCES	Permission is denied to unlink the named semaphore.								
ENAMETOOLONG	The length of <i>name</i> string exceeds <code>PATH_MAX</code> , or a pathname component is longer than <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.								
ENOENT	The named semaphore does not exist.								
ENOSYS	The <code>sem_unlink()</code> function is not supported by the system.								
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
MT-Level	MT-Safe								
SEE ALSO	<code>exec(2)</code> , <code>exit(2)</code> , <code>sem_close(3RT)</code> , <code>sem_open(3RT)</code> , <code>attributes(5)</code>								
NOTES	<p>Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set <code>errno</code> to <code>ENOSYS</code>.</p>								

sem_wait(3RT)

NAME	sem_wait, sem_trywait – acquire or wait for a semaphore										
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <semaphore.h> int sem_wait(sem_t *sem); int sem_trywait(sem_t *sem);</pre>										
DESCRIPTION	<p>The <code>sem_wait()</code> function locks the semaphore referenced by <code>sem</code> by performing a semaphore lock operation on that semaphore. If the semaphore value is currently zero, then the calling thread will not return from the call to <code>sem_wait()</code> until it either locks the semaphore or the call is interrupted by a signal. The <code>sem_trywait()</code> function locks the semaphore referenced by <code>sem</code> only if the semaphore is currently not locked; that is, if the semaphore value is currently positive. Otherwise, it does not lock the semaphore.</p> <p>Upon successful return, the state of the semaphore is locked and remains locked until the <code>sem_post(3RT)</code> function is executed and returns successfully.</p> <p>The <code>sem_wait()</code> function is interruptible by the delivery of a signal.</p>										
RETURN VALUES	The <code>sem_wait()</code> and <code>sem_trywait()</code> functions return 0 if the calling process successfully performed the semaphore lock operation on the semaphore designated by <code>sem</code> . If the call was unsuccessful, the state of the semaphore is unchanged, and the function returns -1 and sets <code>errno</code> to indicate the error.										
ERRORS	<p>The <code>sem_wait()</code> and <code>sem_trywait()</code> functions will fail if:</p> <table><tr><td>EINVAL</td><td>The <code>sem</code> function does not refer to a valid semaphore.</td></tr><tr><td>ENOSYS</td><td>The <code>sem_wait()</code> and <code>sem_trywait()</code> functions are not supported by the system.</td></tr></table> <p>The <code>sem_trywait()</code> function will fail if:</p> <table><tr><td>EAGAIN</td><td>The semaphore was already locked, so it cannot be immediately locked by the <code>sem_trywait()</code> operation.</td></tr></table> <p>The <code>sem_wait()</code> and <code>sem_trywait()</code> functions may fail if:</p> <table><tr><td>EDEADLK</td><td>A deadlock condition was detected; that is, two separate processes are waiting for an available resource to be released via a semaphore "held" by the other process.</td></tr><tr><td>EINTR</td><td>A signal interrupted this function.</td></tr></table>	EINVAL	The <code>sem</code> function does not refer to a valid semaphore.	ENOSYS	The <code>sem_wait()</code> and <code>sem_trywait()</code> functions are not supported by the system.	EAGAIN	The semaphore was already locked, so it cannot be immediately locked by the <code>sem_trywait()</code> operation.	EDEADLK	A deadlock condition was detected; that is, two separate processes are waiting for an available resource to be released via a semaphore "held" by the other process.	EINTR	A signal interrupted this function.
EINVAL	The <code>sem</code> function does not refer to a valid semaphore.										
ENOSYS	The <code>sem_wait()</code> and <code>sem_trywait()</code> functions are not supported by the system.										
EAGAIN	The semaphore was already locked, so it cannot be immediately locked by the <code>sem_trywait()</code> operation.										
EDEADLK	A deadlock condition was detected; that is, two separate processes are waiting for an available resource to be released via a semaphore "held" by the other process.										
EINTR	A signal interrupted this function.										
USAGE	Realtime applications may encounter priority inversion when using semaphores. The problem occurs when a high priority thread "locks" (that is, waits on) a semaphore that is about to be "unlocked" (that is, posted) by a low priority thread, but the low priority thread is preempted by a medium priority thread. This scenario leads to priority inversion; a high priority thread is blocked by lower priority threads for an unlimited period of time. During system design, realtime programmers must take into										

account the possibility of this kind of priority inversion. They can deal with it in a number of ways, such as by having critical sections that are guarded by semaphores execute at a high priority, so that a thread cannot be preempted while executing in its critical section.

EXAMPLES **EXAMPLE 1** The customer waiting-line in a bank may be analogous to the synchronization scheme of a semaphore utilizing `sem_wait()` and `sem_trywait()`:

```
/* cc [flag ...] file ... -lrt -pthread [library ...] */

#include <errno.h>
#define TELLERS 10
sem_t bank_line;      /* semaphore */
int banking_hours(), deposit_withdrawal;
void *customer(), do_business(), skip_banking_today();
thread_t tid;
. . .

sem_init(&bank_line, TRUE, TELLERS); /* 10 tellers available */
while(banking_hours())
    thr_create(NULL, NULL, customer, (void *)deposit_withdrawal,
              THREAD_NEW_LWP, &tid);
. . .

void *
customer(deposit_withdrawal)
void *deposit_withdrawal;
{
    int this_customer, in_a_hurry = 50;
    this_customer = rand() % 100;
    if (this_customer == in_a_hurry) {
        if (sem_trywait(&bank_line) != 0)
            if (errno == EAGAIN) { /* no teller available */
                skip_banking_today(this_customer);
                return;
            } /*else go immediately to available teller
              & decrement bank_line*/
    }
    else
        sem_wait(&bank_line); /* wait for next teller,
                              then proceed, and decrement bank_line */
    do_business((int *)deposit_withdrawal);
    sem_getvalue(&bank_line, &num_tellers);
    sem_post(&bank_line); /* increment bank_line;
                          this_customer's teller is now available */
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

sem_wait(3RT)

SEE ALSO | sem_post(3RT), attributes(5)

NAME	shm_open – open a shared memory object								
SYNOPSIS	<pre>cc [<i>flag...</i>] <i>file...</i> -lrt [<i>library...</i>] #include <sys/mman.h> int shm_open(const char *name, int oflag, mode_t mode);</pre>								
DESCRIPTION	<p>The shm_open() function establishes a connection between a shared memory object and a file descriptor. It creates an open file description that refers to the shared memory object and a file descriptor that refers to that open file description. The file descriptor is used by other functions to refer to that shared memory object. The <i>name</i> argument points to a string naming a shared memory object. It is unspecified whether the name appears in the file system and is visible to other functions that take pathnames as arguments. The <i>name</i> argument conforms to the construction rules for a pathname. The first character of <i>name</i> must be a slash (/) character and the remaining characters of <i>name</i> cannot include any slash characters. For maximum portability, <i>name</i> should include no more than 14 characters, but this limit is not enforced.</p> <p>If successful, shm_open() returns a file descriptor for the shared memory object that is the lowest numbered file descriptor not currently open for that process. The open file description is new, and therefore the file descriptor does not share it with any other processes. It is unspecified whether the file offset is set. The FD_CLOEXEC file descriptor flag associated with the new file descriptor is set.</p> <p>The file status flags and file access modes of the open file description are according to the value of <i>oflag</i>. The <i>oflag</i> argument is the bitwise inclusive OR of the following flags defined in the header <fcntl.h>. Applications specify exactly one of the first two values (access modes) below in the value of <i>oflag</i>:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">O_RDONLY</td> <td>Open for read access only.</td> </tr> <tr> <td>O_RDWR</td> <td>Open for read or write access.</td> </tr> </table> <p>Any combination of the remaining flags may be specified in the value of <i>oflag</i>:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">O_CREAT</td> <td>If the shared memory object exists, this flag has no effect, except as noted under O_EXCL below. Otherwise the shared memory object is created; the user ID of the shared memory object will be set to the effective user ID of the process; the group ID of the shared memory object will be set to a system default group ID or to the effective group ID of the process. The permission bits of the shared memory object will be set to the value of the <i>mode</i> argument except those set in the file mode creation mask of the process. When bits in <i>mode</i> other than the file permission bits are set, the effect is unspecified. The <i>mode</i> argument does not affect whether the shared memory object is opened for reading, for writing, or for both. The shared memory object has a size of zero.</td> </tr> <tr> <td>O_EXCL</td> <td>If O_EXCL and O_CREAT are set, shm_open() fails if the shared memory object exists. The check for the existence of the shared memory object and the creation of the object if it does not exist is</td> </tr> </table>	O_RDONLY	Open for read access only.	O_RDWR	Open for read or write access.	O_CREAT	If the shared memory object exists, this flag has no effect, except as noted under O_EXCL below. Otherwise the shared memory object is created; the user ID of the shared memory object will be set to the effective user ID of the process; the group ID of the shared memory object will be set to a system default group ID or to the effective group ID of the process. The permission bits of the shared memory object will be set to the value of the <i>mode</i> argument except those set in the file mode creation mask of the process. When bits in <i>mode</i> other than the file permission bits are set, the effect is unspecified. The <i>mode</i> argument does not affect whether the shared memory object is opened for reading, for writing, or for both. The shared memory object has a size of zero.	O_EXCL	If O_EXCL and O_CREAT are set, shm_open() fails if the shared memory object exists. The check for the existence of the shared memory object and the creation of the object if it does not exist is
O_RDONLY	Open for read access only.								
O_RDWR	Open for read or write access.								
O_CREAT	If the shared memory object exists, this flag has no effect, except as noted under O_EXCL below. Otherwise the shared memory object is created; the user ID of the shared memory object will be set to the effective user ID of the process; the group ID of the shared memory object will be set to a system default group ID or to the effective group ID of the process. The permission bits of the shared memory object will be set to the value of the <i>mode</i> argument except those set in the file mode creation mask of the process. When bits in <i>mode</i> other than the file permission bits are set, the effect is unspecified. The <i>mode</i> argument does not affect whether the shared memory object is opened for reading, for writing, or for both. The shared memory object has a size of zero.								
O_EXCL	If O_EXCL and O_CREAT are set, shm_open() fails if the shared memory object exists. The check for the existence of the shared memory object and the creation of the object if it does not exist is								

shm_open(3RT)

atomic with respect to other processes executing `shm_open()` naming the same shared memory object with `O_EXCL` and `O_CREAT` set. If `O_EXCL` is set and `O_CREAT` is not set, the result is undefined.

`O_TRUNC` If the shared memory object exists, and it is successfully opened `O_RDWR`, the object will be truncated to zero length and the mode and owner will be unchanged by this function call. The result of using `O_TRUNC` with `O_RDONLY` is undefined.

When a shared memory object is created, the state of the shared memory object, including all data associated with the shared memory object, persists until the shared memory object is unlinked and all other references are gone. It is unspecified whether the name and shared memory object state remain valid after a system reboot.

RETURN VALUES Upon successful completion, the `shm_open()` function returns a non-negative integer representing the lowest numbered unused file descriptor. Otherwise, it returns `-1` and sets `errno` to indicate the error condition.

ERRORS The `shm_open()` function will fail if:

<code>EACCES</code>	The shared memory object exists and the permissions specified by <i>oflag</i> are denied, or the shared memory object does not exist and permission to create the shared memory object is denied, or <code>O_TRUNC</code> is specified and write permission is denied.
<code>EEXIST</code>	<code>O_CREAT</code> and <code>O_EXCL</code> are set and the named shared memory object already exists.
<code>EINTR</code>	The <code>shm_open()</code> operation was interrupted by a signal.
<code>EINVAL</code>	The <code>shm_open()</code> operation is not supported for the given name.
<code>EMFILE</code>	Too many file descriptors are currently in use by this process.
<code>ENAMETOOLONG</code>	The length of the <i>name</i> string exceeds <code>PATH_MAX</code> , or a pathname component is longer than <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
<code>ENFILE</code>	Too many shared memory objects are currently open in the system.
<code>ENOENT</code>	<code>O_CREAT</code> is not set and the named shared memory object does not exist.
<code>ENOSPC</code>	There is insufficient space for the creation of the new shared memory object.

ENOSYS The shm_open () function is not supported by the system.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO close(2), dup(2), exec(2), fcntl(2), mmap(2), umask(2), shm_unlink(3RT), sysconf(3C), attributes(5), fcntl(3HEAD)

NOTES Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set errno to ENOSYS.

shm_unlink(3RT)

NAME	shm_unlink – remove a shared memory object				
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <sys/mman.h> int shm_unlink(const char *name);</pre>				
DESCRIPTION	The shm_unlink() function removes the name of the shared memory object named by the string pointed to by <i>name</i> . If one or more references to the shared memory object exists when the object is unlinked, the name is removed before shm_unlink() returns, but the removal of the memory object contents will be postponed until all open and mapped references to the shared memory object have been removed.				
RETURN VALUES	Upon successful completion, shm_unlink() returns 0. Otherwise it returns -1 and sets errno to indicate the error condition, and the named shared memory object is not affected by this function call.				
ERRORS	The shm_unlink() function will fail if: EACCES Permission is denied to unlink the named shared memory object. ENAMETOOLONG The length of the <i>name</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX while _POSIX_NO_TRUNC is in effect. ENOENT The named shared memory object does not exist. ENOSYS The shm_unlink() function is not supported by the system.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	close(2), mmap(2), mlock(3C), shm_open(3RT), attributes(5)				
NOTES	Solaris 2.6 was the first release to support the Asynchronous Input and Output option. Prior to this release, this function always returned -1 and set errno to ENOSYS.				

NAME	sigqueue – queue a signal to a process										
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <sys/types.h> #include <signal.h> int sigqueue(pid_t pid, int signo, const union signal value);</pre>										
DESCRIPTION	<p>The <code>sigqueue()</code> function causes the signal specified by <code>signo</code> to be sent with the value specified by <code>value</code> to the process specified by <code>pid</code>. If <code>signo</code> is 0 (the null signal), error checking is performed but no signal is actually sent. The null signal can be used to check the validity of <code>pid</code>.</p> <p>The conditions required for a process to have permission to queue a signal to another process are the same as for the <code>kill(2)</code> function.</p> <p>The <code>sigqueue()</code> function returns immediately. If <code>SA_SIGINFO</code> is set for <code>signo</code> and if the resources were available to queue the signal, the signal is queued and sent to the receiving process. If <code>SA_SIGINFO</code> is not set for <code>signo</code>, then <code>signo</code> is sent at least once to the receiving process; it is unspecified whether <code>value</code> will be sent to the receiving process as a result of this call.</p> <p>If the value of <code>pid</code> causes <code>signo</code> to be generated for the sending process, and if <code>signo</code> is not blocked for the calling thread and if no other thread has <code>signo</code> unblocked or is waiting in a <code>sigwait(2)</code> function for <code>signo</code>, either <code>signo</code> or at least the pending, unblocked signal will be delivered to the calling thread before the <code>sigqueue()</code> function returns. Should any of multiple pending signals in the range <code>SIGRTMIN</code> to <code>SIGRTMAX</code> be selected for delivery, it will be the lowest numbered one. The selection order between realtime and non-realtime signals, or between multiple pending non-realtime signals, is unspecified.</p>										
RETURN VALUES	Upon successful completion, the specified signal will have been queued, and the <code>sigqueue()</code> function returns 0. Otherwise, the function returns -1 and sets <code>errno</code> to indicate the error.										
ERRORS	<p>The <code>sigqueue()</code> function will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EAGAIN</td> <td>No resources are available to queue the signal. The process has already queued <code>SIGQUEUE_MAX</code> signals that are still pending at the receiver(s), or a system wide resource limit has been exceeded.</td> </tr> <tr> <td style="vertical-align: top;">EINVAL</td> <td>The value of <code>signo</code> is an invalid or unsupported signal number.</td> </tr> <tr> <td style="vertical-align: top;">ENOSYS</td> <td>The <code>sigqueue()</code> function is not supported by the system.</td> </tr> <tr> <td style="vertical-align: top;">EPERM</td> <td>The process does not have the appropriate privilege to send the signal to the receiving process.</td> </tr> <tr> <td style="vertical-align: top;">ESRCH</td> <td>The process <code>pid</code> does not exist.</td> </tr> </table>	EAGAIN	No resources are available to queue the signal. The process has already queued <code>SIGQUEUE_MAX</code> signals that are still pending at the receiver(s), or a system wide resource limit has been exceeded.	EINVAL	The value of <code>signo</code> is an invalid or unsupported signal number.	ENOSYS	The <code>sigqueue()</code> function is not supported by the system.	EPERM	The process does not have the appropriate privilege to send the signal to the receiving process.	ESRCH	The process <code>pid</code> does not exist.
EAGAIN	No resources are available to queue the signal. The process has already queued <code>SIGQUEUE_MAX</code> signals that are still pending at the receiver(s), or a system wide resource limit has been exceeded.										
EINVAL	The value of <code>signo</code> is an invalid or unsupported signal number.										
ENOSYS	The <code>sigqueue()</code> function is not supported by the system.										
EPERM	The process does not have the appropriate privilege to send the signal to the receiving process.										
ESRCH	The process <code>pid</code> does not exist.										
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:										

sigqueue(3RT)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO kill(2), sigwaitinfo(3RT), attributes(5), siginfo(3HEAD), signal(3HEAD)

NAME	sigwaitinfo, sigtimedwait – wait for queued signals
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <signal.h> int sigwaitinfo(const sigset_t *set, siginfo_t *info); int sigtimedwait(const sigset_t *set, siginfo_t *info, const struct timespec *timeout);</pre>
DESCRIPTION	<p>The sigwaitinfo() function selects the pending signal from the set specified by set. Should any of multiple pending signals in the range SIGRTMIN to SIGRTMAX be selected, it will be the lowest numbered one. The selection order between realtime and non-realtime signals, or between multiple pending non-realtime signals, is unspecified. If no signal in set is pending at the time of the call, the calling thread is suspended until one or more signals in set become pending or until it is interrupted by an unblocked, caught signal.</p> <p>The sigwaitinfo() function behaves the same as the sigwait(2) function if the info argument is NULL. If the info argument is non-NULL, the sigwaitinfo() function behaves the same as sigwait(2), except that the selected signal number is stored in the si_signo member, and the cause of the signal is stored in the si_code member. If any value is queued to the selected signal, the first such queued value is dequeued and, if the info argument is non-NULL, the value is stored in the si_value member of info. The system resource used to queue the signal will be released and made available to queue other signals. If no value is queued, the content of the si_value member is undefined. If no further signals are queued for the selected signal, the pending indication for that signal will be reset. If the value of the si_code member is SI_NOINFO, only the si_signo member of siginfo_t is meaningful, and the value of all other members is unspecified.</p> <p>The sigtimedwait() function behaves the same as sigwaitinfo() except that if none of the signals specified by set are pending, sigtimedwait() waits for the time interval specified in the timespec structure referenced by timeout. If the timespec structure pointed to by timeout is zero-valued and if none of the signals specified by set are pending, then sigtimedwait() returns immediately with an error. If timeout is the NULL pointer, the behavior is unspecified.</p> <p>If, while sigwaitinfo() or sigtimedwait() is waiting, a signal occurs which is eligible for delivery (that is, not blocked by the process signal mask), that signal is handled asynchronously and the wait is interrupted.</p>
RETURN VALUES	Upon successful completion (that is, one of the signals specified by set is pending or is generated) sigwaitinfo() and sigtimedwait() will return the selected signal number. Otherwise, the function returns -1 and sets errno to indicate the error.
ERRORS	<p>The sigwaitinfo() and sigtimedwait() functions will fail if:</p> <p>EINTR The wait was interrupted by an unblocked, caught signal.</p>

sigwaitinfo(3RT)

ENOSYS The `sigwaitinfo()` and `sigtimedwait()` functions are not supported.

The `sigtimedwait()` function will also fail if:

EAGAIN No signal specified by `set` was generated within the specified timeout period.

The `sigtimedwait()` function may also fail if:

EINVAL The `timeout` argument specified a `tv_nsec` value less than zero or greater than or equal to 1000 million. The system only checks for this error if no signal is pending in `set` and it is necessary to wait.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Safe

SEE ALSO `time(2)`, `sigqueue(3RT)`, `attributes(5)`, `siginfo(3HEAD)`, `signal(3HEAD)`, `time(3HEAD)`

NAME | td_init – performs initialization for libthread_db library of interfaces

SYNOPSIS | `cc [flag ...] file ... -lthread_db [library ...]`

```
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_init();
```

DESCRIPTION | td_init() is the global initialization function for the libthread_db() library of interfaces. It must be called exactly once by any process using the libthread_db() library before any other libthread_db function can be called.

RETURN VALUES | TD_OK The libthread_db() library of interfaces successfully initialized.

 | TD_ERR Initialization failed.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO | libthread_db(3THR), libthread_db(3LIB), attributes(5)

td_log(3THR)

NAME | td_log – placeholder for future logging functionality

SYNOPSIS | `cc [flag ...] file ... -lthread_db [library ...]`

```
#include <proc_service.h>
#include <thread_db.h>

void td_log();
```

DESCRIPTION | This function presently does nothing; it is merely a placeholder for future logging functionality in libthread_db(3THR).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO | libthread(3LIB), libthread_db(3LIB), threads(3THR), attributes(5)

td_sync_get_info(3THR)

NAME	td_sync_get_info, td_ta_sync_tracking_enable, td_sync_get_stats, td_sync_setstate, td_sync_waiters – operations on a synchronization object in libthread_db
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i> ... -lthread_db [<i>library</i> ...] #include <proc_service.h> #include <thread_db.h> td_err_e td_sync_get_info(const td_synchandle_t *<i>sh_p</i>, td_syncinfo_t *<i>si_p</i>); td_err_e td_ta_sync_tracking_enable(const td_thragent_t *<i>ta_p</i>, int <i>on_off</i>); td_err_e td_sync_get_stats(const td_synchandle_t *<i>sh_p</i>, td_syncstats_t *<i>ss_p</i>); td_err_e td_sync_setstate(const td_synchandle_t *<i>sh_p</i>); typedef int td_thr_iter_f(const td_thrhandle_t *<i>th_p</i>, void *<i>cb_data_p</i>); td_err_e td_sync_waiters(const td_synchandle_t *<i>sh_p</i>, td_thr_iter_f *<i>cb</i>, void *<i>cb_data_p</i>);</pre>
DESCRIPTION	<p>Synchronization objects include mutexes, condition variables, semaphores, and reader-writer locks. In the same way that thread operations use a thread handle of type <code>td_thrhandle_t</code>, operations on synchronization objects use a synchronization object handle of type <code>td_synchandle_t</code>.</p> <p>The controlling process obtains synchronization object handles either by calling the function <code>td_ta_sync_iter()</code> to obtain handles for all synchronization objects of the target process that are known to the <code>libthread_db</code> library of interfaces, or by mapping the address of a synchronization object in the address space of the target process to a handle by calling <code>td_ta_map_addr2sync(3THR)</code>.</p> <p>Not all synchronization objects that a process uses may be known to the <code>libthread_db</code> library and returned by <code>td_ta_sync_iter(3THR)</code>. A synchronization object is known to <code>libthread_db</code> only if it has been the target of a synchronization primitive in the process (such as <code>mutex_lock()</code>, described on the <code>mutex_init(3THR)</code> manual page) after <code>td_ta_new(3THR)</code> has been called to attach to the process and <code>td_ta_sync_tracking_enable()</code> has been called to enable synchronization object tracking.</p> <p>The <code>td_ta_sync_tracking_enable()</code> function turns synchronization object tracking on or off for the process identified by <i>ta_p</i>, depending on whether <i>on_off</i> is 0 (off) or non-zero (on).</p> <p>The <code>td_sync_get_info()</code> function fills in the <code>td_syncinfo_t</code> structure *<i>si_p</i> with values for the synchronization object identified by <i>sh_p</i>. The <code>td_syncinfo_t</code> structure contains the following fields:</p>

td_sync_get_info(3THR)

<code>td_thragent_t *si_ta_p</code>	The internal process handle identifying the target process through which this synchronization object handle was obtained. Synchronization objects may be process-private or process-shared. In the latter case, the same synchronization object may have multiple handles, one for each target process's "view" of the synchronization object.
<code>psaddr_t si_sv_addr</code>	The address of the synchronization object in this target process's address space.
<code>td_sync_type_e si_type</code>	The type of the synchronization variable: mutex, condition variable, semaphore, or readers-writer lock.
<code>int si_shared_type</code>	If the <code>USYNC_PROCESS</code> or <code>USYNC_PROCESS_ROBUST</code> bit is set in <code>si_shared_type</code> , this synchronization object is process-shared, otherwise it is process-private.
<code>td_sync_flags_t si_flags</code>	Flags dependent on the type of the synchronization object.
<code>int si_state.sema_count</code>	Semaphores only. The current value of the semaphore
<code>int si_state.nreaders</code>	Readers-writer locks only. The number of readers currently holding the lock, or -1, if a writer is currently holding the lock.
<code>int si_state.mutex_locked</code>	For mutexes only. Non-zero if and only if the mutex is currently locked.
<code>int si_size</code>	The size of the synchronization object.
<code>uint8_t si_has_waiters</code>	Non-zero if and only if at least one thread is blocked on this synchronization object.
<code>uint8_t si_is_wlocked</code>	For reader-writer locks only. The value is non-zero if and only if this lock is held by a writer.
<code>uint8_t si_rcount</code>	<code>PTHREAD_MUTEX_RECURSIVE</code> mutexes only. If the mutex is held, the recursion count.
<code>uint8_t si_prioceiling</code>	<code>PTHREAD_PRIO_PROTECT</code> protocol mutexes only. The priority ceiling.

td_sync_get_info(3THR)

td_thrhandle_t *si_owner* Mutexes and readers-writer locks only. This is the thread holding the mutex, or the write lock, if this is a reader-writer lock. The value is NULL if no one holds the mutex or write-lock.

pid_t *si_ownerpid* Mutexes only. For a locked process-shared mutex, this is the process-ID of the process containing the owning thread.

The `td_sync_get_stats()` function fills in the `td_syncstats_t` structure `*ss_p` with values for the synchronization object identified by `sh_p`. The `td_syncstats_t` structure contains an embedded `td_syncinfo_t` structure that is filled in as described above for `td_sync_get_info()`. In addition, usage statistics gathered since `td_ta_sync_tracking_enable()` was called to enable synchronization object tracking are returned in the `ss_un.mutex`, `ss_un.cond`, `ss_un.rwlock`, or `ss_un.sema` members of the `td_syncstats_t` structure, depending on the type of the synchronization object.

The `td_sync_setstate` function modifies the state of synchronization object `si_p`, depending on the synchronization object type. For mutexes, `td_sync_setstate` is unlocked if the value is 0. Otherwise it is locked. For semaphores, the semaphore's count is set to the value. For reader-writer locks, the reader count set to the value if value is >0. The count is set to write-locked if value is -1. It is set to unlocked if the value is 0. Setting the state of a synchronization object from a `libthread_db` interface may cause the synchronization object's semantics to be violated from the point of view of the threads in the target process. For example, if a thread holds a mutex, and `td_sync_setstate` is used to set the mutex to unlocked, then a different thread will also be able to subsequently acquire the same mutex.

The `td_sync_waiters` function iterates over the set of thread handles of threads blocked on `sh_p`. The callback function `cb` is called once for each such thread handle, and is passed the thread handle and `cb_data_p`. If the callback function returns a non-zero value, iteration is terminated early. See `td_ta_thr_iter(3THR)`.

RETURN VALUES

TD_OK The call returned successfully.

TD_BADTH An invalid thread handle was passed in.

TD_DBERR A call to one of the imported interface routines failed.

TD_ERR A `libthread_db`-internal error occurred.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

td_sync_get_info(3THR)

SEE ALSO | libthread_db(3LIB), libthread_db(3THR), mutex_init(3THR),
td_ta_map_addr2sync(3THR), td_ta_sync_iter(3THR),
td_ta_thr_iter(3THR), attributes(5)

NAME	td_ta_enable_stats, td_ta_reset_stats, td_ta_get_stats – collect target process statistics for libthread_db
SYNOPSIS	<pre>cc [flag ...] file ... -lthread_db [library ...] #include <proc_service.h> #include <thread_db.h> td_err_e td_ta_enable_stats(const td_thragent_t *ta_p, int on_off); td_err_e _statstd_ta_reset(const td_thragent_t *ta_p); td_err_e td_ta_get_stats(const td_thragent_t *ta_p, td_ta_stats_t *tstats);</pre>
DESCRIPTION	<p>The controlling process may request the collection of certain statistics about a target process. Statistics gathering is disabled by default; however, each target process has a <code>td_ta_stats_t</code> structure that contains up to date values when statistic gathering is enabled. <code>td_ta_enable_stats()</code> turns statistics gathering on or off for the process identified by <code>ta_p</code> depending on whether or not <code>on_off</code> is non-zero. When statistics gathering is turned on, all statistics are implicitly reset as though <code>td_ta_reset_stats()</code> had been called. Statistics are not reset when statistics gathering is turned off. Except for <code>nthreads</code> and <code>r_concurrency</code>, the values do not change further, but they remain available for inspection by way of <code>td_ta_get_stats()</code>. <code>td_ta_reset_stats()</code> resets all counters in the <code>td_ta_stats_t</code> structure to zero for the target process. <code>td_ta_get_stats()</code> returns the <code>td_ta_stats_t</code> structure for the process in <code>*stats_t</code>. The <code>td_ta_stats_t</code> structure is defined as follows:</p> <pre>typedef struct { int nthreads; /* total number of threads in use */ int r_concurrency; /* requested concurrency level */ int nrunnable_num; /* numerator of avg. runnable threads */ int nrunnable_den; /* denominator of avg. runnable threads */ int a_concurrency_num; /* numerator, avg. achieved concurrency */ int a_concurrency_den; /* denominator, avg. achieved concurrency */ int nlwps_num; /* numerator, average number of LWPs in use */ int nlwps_den; /* denominator, avg. number of LWPs in use */ int nidle_num; /* numerator, avg. number of idling LWPs */ int nidle_den; /* denominator, avg. number of idling LWPs */ } td_ta_stats_t;</pre> <p><code>nthreads</code> is the number of threads that are currently part of the target process. <code>r_concurrency</code> is the current requested concurrency level, such as would be returned by <code>thr_setconcurrency(3THR)</code>. The remaining fields are averages over time, each expressed as a fraction with an integral numerator and denominator. <code>nrunnable</code> is the average number of runnable threads. <code>a_concurrency</code> is the average achieved concurrency, the number of actually running threads. <code>a_concurrency</code> is less than or equal to <code>nrunnable</code>. <code>nlwps</code> is the average number of lightweight processes (LWPs) participating in this process. It must be greater than or equal to <code>a_concurrency</code>, as every running thread is assigned to an LWP, but there may at times be additional idling LWPs with no thread assigned to them. <code>nidle</code> is the average number of idle LWPs.</p>

td_ta_enable_stats(3THR)

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTA	An invalid internal process handle was passed in.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	Something else went wrong.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO `libthread_db(3THR)`, `thr_getconcurrency(3THR)`, `libthread_db(3LIB)`, `attributes(5)`

td_ta_event_addr(3THR)

NAME	td_ta_event_addr, td_thr_event_enable, td_ta_set_event, td_thr_set_event, td_ta_clear_event, td_thr_clear_event, td_ta_event_getmsg, td_thr_event_getmsg, td_event_emptyset, td_event_fillset, td_event_addset, td_event_delset, td_eventismember, td_eventisempty – thread events in libthread_db
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i> ... -lthread_db [<i>library</i> ...] #include <proc_service.h> #include <thread_db.h> td_err_e td_ta_event_addr(const td_thragent_t *<i>ta_p</i>, u_long event, td_notify_t *<i>notify_p</i>); td_err_e td_thr_event_enable(const td_thrhandle_t *<i>th_p</i>, int <i>on_off</i>); td_err_e td_thr_set_event(const td_thrhandle_t *<i>th_p</i>, td_thr_events_t *<i>events</i>); td_err_e td_ta_set_event(const td_thragent_t *<i>ta_p</i>, td_thr_events_t *<i>events</i>); td_err_e td_thr_clear_event(const td_thrhandle_t *<i>th_p</i>, td_thr_events_t *<i>events</i>); td_err_e td_ta_clear_event(const td_thragent_t *<i>ta_p</i>, td_thr_events_t *<i>events</i>); td_err_e td_thr_event_getmsg(const td_thrhandle_t *<i>th_p</i>, td_event_msg_t *<i>msg</i>); td_err_e td_ta_event_getmsg(const td_thragent_t *<i>ta_p</i>, td_event_msg_t *<i>msg</i>); void td_event_emptyset(td_thr_events_t *); void td_event_fillset(td_thr_events_t *); void td_event_addset(td_thr_events_t *, td_thr_events_e <i>n</i>); void td_event_delset(td_thr_events_t *, td_thr_events_e <i>n</i>); void td_eventismember(td_thr_events_t *, td_thr_events_e <i>n</i>); void td_eventisempty(td_thr_events_t*);</pre>
DESCRIPTION	<p>These routines comprise the thread event facility for libthread_db(3THR). This facility allows the controlling process to be notified when certain thread-related events occur in a target process and to retrieve information associated with these events. An event consists of an event type, and optionally, some associated event data, depending on the event type. See the section titled "Event Set Manipulation Macros" that follows.</p> <p>The event type and the associated event data, if any, constitute an "event message." "Reporting an event" means delivering an event message to the controlling process by way of libthread_db.</p>

td_ta_event_addr(3THR)

Several flags can control event reporting, both a per-thread and per event basis. Event reporting may further be enabled or disabled for a thread. There is not only a per-thread event mask that specifies which event types should be reported for that thread, but there is also a global event mask that applies to all threads.

An event is reported, if and only if, the executing thread has event reporting enabled, and either the event type is enabled in the executing thread's event mask, or the event type is enabled in the global event mask.

Each thread has associated with it an event buffer in which it stores the most recent event message it has generated, the type of the most recent event that it reported, and, depending on the event type, some additional information related to that event. See the section titled "Event Set Manipulation Macros" for a description of the `td_thr_events_e` and `td_event_msg_t` types and a list of the event types and the values reported with them. The thread handle, type `td_thrhandle_t`, the event type, and the possible value, together constitute an event message. Each thread's event buffer holds at most one event message.

Each event type has an event reporting address associated with it. A thread reports an event by writing the event message into the thread's event buffer and having control reach the event reporting address for that event type.

Typically, the controlling process sets a breakpoint at the event reporting address for one or more event types. When the breakpoint is hit, the controlling process knows that an event of the corresponding type has occurred.

The event types, and the additional information, if any, reported with each event, are:

<code>TD_READY</code>	The thread became ready to execute.
<code>TD_SLEEP</code>	The thread has blocked on a synchronization object.
<code>TD_SWITCHTO</code>	A runnable thread is being assigned to LWP.
<code>TD_SWITCHFROM</code>	A running thread is being removed from its LWP.
<code>TD_LOCK_TRY</code>	A thread is trying to get an unavailable lock.
<code>TD_CATCHSIG</code>	A signal was posted to a thread.
<code>TD_IDLE</code>	An LWP is becoming idle.
<code>TD_CREATE</code>	A thread is being created.
<code>TD_DEATH</code>	A thread has terminated.
<code>TD_PREEMPT</code>	A thread is being preempted.
<code>TD_PRI_INHERIT</code>	A thread is inheriting an elevated priority from another thread.
<code>TD_REAP</code>	A thread is being reaped.
<code>TD_CONCURRENCY</code>	The number of LWPs is changing.

td_ta_event_addr(3THR)

TD_TIMEOUT A condition-variable timed wait expired.

td_ta_event_addr() returns in **notify_p* the event reporting address associated with event type *event*. The controlling process may then set a breakpoint at that address. If a thread hits that breakpoint, it reports an event of type *event*.

td_thr_event_enable() enables or disables event reporting for thread *th_p*. If a thread has event reporting disabled, it will not report any events. Threads are started with event reporting disabled. Event reporting is enabled if *on_off* is non-zero; otherwise, it is disabled. To find out whether or not event reporting is enabled on a thread, call *td_thr_getinfo()* for the thread and examine the *ti_traceme* field of the *td_thrinfo_t* structure it returns.

td_thr_set_event() and *td_thr_clear_event()* set and clear, respectively, a set of event types in the event mask associated with the thread *th_p*. To inspect a thread's event mask, call *td_thr_getinfo()* for the thread, and examine the *ti_events* field of the *td_thrinfo_t* structure it returns.

td_ta_set_event() and *td_ta_clear_event()* are just like *td_thr_set_event()* and *td_thr_clear_event()*, respectively, except that the target process's global event mask is modified. There is no provision for inspecting the value of a target process's global event mask.

td_thr_event_getmsg() returns in **msg* the event message associated with thread **th_p*. Reading a thread's event message consumes the message, emptying the thread's event buffer. As noted above, each thread's event buffer holds at most one event message; if a thread reports a second event before the first event message has been read, the second event message overwrites the first.

td_ta_event_getmsg() is just like *td_thr_event_getmsg()*, except that it is passed a process handle rather than a thread handle. It selects some thread that has an event message buffered, and it returns that thread's message. The thread selected is undefined, except that as long as at least one thread has an event message buffered, it will return an event message from some such thread.

Event Set Manipulation Macros

Several macros are provided for manipulating event sets of type *td_thr_events_t*:

<i>td_event_emptyset</i>	Sets its argument to the NULL event set.
<i>td_event_fillset</i>	Sets its argument to the set of all events.
<i>td_event_addset</i>	Adds a specific event type to an event set.
<i>td_event_delset</i>	Deletes a specific event type from an event set.
<i>td_eventismember</i>	Tests whether a specific event type is a member of an event set.
<i>td_eventisempty</i>	Tests whether an event set is the NULL set.

RETURN VALUES

The following values may be returned for all thread event routines:

td_ta_event_addr(3THR)

TD_OK	The call returned successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_BADTA	An invalid internal process handle was passed in.
TD_BADPH	There is a NULL external process handle associated with this internal process handle.
TD_DBERR	A call to one of the imported interface routines failed.
TD_NOMSG	No event message was available to return to <code>td_thr_event_getmsg()</code> or <code>td_ta_event_getmsg()</code> .
TD_ERR	Some other parameter error occurred, or a <code>libthread_db()</code> internal error occurred.

The following value may be returned for `td_thr_event_enable()`, `td_thr_set_event()`, and `td_thr_clear_event()` only:

TD_NOCAPAB	The agent thread in the target process has not completed initialization, so this operation cannot be performed. The operation can be performed after the target process has been allowed to make some forward progress. See also <code>libthread_db(3THR)</code> .
------------	--

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libthread_db(3THR)`, `libthread_db(3LIB)`, `attributes(5)`

td_ta_get_nthreads(3THR)

NAME | td_ta_get_nthreads – gets the total number of threads in a process for libthread_db

SYNOPSIS |

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

|

```
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_ta_get_nthreads(const td_thragent_t *ta_p, int
    *nthread_p);
```

DESCRIPTION | The td_ta_get_nthreads() function returns the total number of threads in process *ta_p*, including any system threads. System threads are those created by libthread or libthread_db on its own behalf. The number of threads is written into **nthread_p*.

RETURN VALUES | TD_OK The call completed successfully.

| TD_BADTA An invalid internal process handle was passed in.

| TD_BADPH There is a NULL external process handle associated with this internal process handle.

| TD_DBERR A call to one of the imported interface routines failed.

| TD_ERR *nthread_p* was NULL, or a libthread_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | libthread(3LIB), libthread_db(3THR), libthread_db(3LIB), threads(3THR), attributes(5)

td_ta_map_addr2sync(3THR)

NAME td_ta_map_addr2sync – get a synchronization object handle from a synchronization object's address

SYNOPSIS

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

```
#include <proc_service.h>
#include <thread_db.h>

td_ta_map_addr2sync(const td_thragent_t *ta_p, psaddr_t
    addr, td_synchandle_t *sh_p);
```

DESCRIPTION td_ta_map_addr2sync() produces the synchronization object handle of type td_synchandle_t that corresponds to the address of the synchronization object (mutex, semaphore, condition variable, or reader/writer lock). Some effort is made to validate *addr* and verify that it does indeed point at a synchronization object. The handle is returned in **sh_p*.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTA	An invalid internal process handle was passed in.
TD_BADPH	There is a NULL external process handle associated with this internal process handle.
TD_BADSH	<i>sh_p</i> is NULL, or <i>addr</i> does not appear to point to a valid synchronization object.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	<i>addr</i> is NULL, or a libthread_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO libthread_db(3THR), libthread_db(3LIB), attributes(5)

NAME td_ta_map_id2thr, td_ta_map_lwp2thr – convert a thread id or LWP id to a thread handle

SYNOPSIS

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

#include <proc_service.h>
#include <thread_db.h>

td_ta_map_id2thr(const td_thragent_t *ta_p, thread_t
    tid, td_thrhandle_t *th_p);

td_ta_map_lwp2thr(const td_thragent_t *ta_p, lwpid_t
    lwpid, td_thrhandle_t *th_p);
```

DESCRIPTION

td_ta_map_id2thr() produces the td_thrhandle_t thread handle that corresponds to a particular thread id, as returned by thr_create(3THR) or thr_self(3THR). The thread handle is returned in *th_p*.

td_ta_map_lwp2thr() produces the td_thrhandle_t thread handle for the thread that is currently executing on the light weight process (LWP) and has an id of lwpid.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTA	An invalid internal process handle was passed in.
TD_BADPH	There is a NULL external process handle associated with this internal process handle.
TD_DBERR	A call to one of the imported interface routines failed.
TD_NOTHR	Either there is no thread with the given thread id (td_ta_map_id2thr) or no thread is currently executing on the given LWP (td_ta_map_lwp2thr).
TD_ERR	The call did not complete successfully.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO libthread_db(3THR), thr_create(3THR), thr_self(3THR), libthread_db(3LIB), attributes(5)

td_ta_new(3THR)

NAME	td_ta_new, td_ta_delete, td_ta_get_ph – allocate and deallocate process handles for libthread_db												
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i> ... -lthread_db [<i>library</i> ...] #include <proc_service.h> #include <thread_db.h> td_err_e td_ta_new(const struct ps_prochandle *<i>ph_p</i>, td_thragent_t **<i>ta_pp</i>); td_err_e td_ta_delete(const td_thragent_t *<i>ta_p</i>); td_err_e td_ta_get_ph(const td_thragent_t *<i>ta_p</i>, struct ps_prochandle **<i>ph_pp</i>);</pre>												
DESCRIPTION	<p>td_ta_new() registers a target process with libthread_db and allocates an internal process handle of type td_thragent_t for this target process. Subsequent calls to libthread_db can use this handle to refer to this target process.</p> <p>There are actually two process handles, an internal process handle assigned by libthread_db and an external process handle assigned by the libthread_db client. There is a one-to-one correspondence between the two handles. When the client calls a libthread_db routine, it uses the internal process handle. When libthread_db calls one of the client-provided routines listed in proc_service(3PROC), it uses the external process handle.</p> <p><i>ph</i> is the external process handle that libthread_db should use to identify this target process to the controlling process when it calls routines in the imported interface.</p> <p>If this call is successful, the value of the newly allocated td_thragent_t handle is returned in *<i>ta_pp</i>. td_ta_delete() deregisters a target process with libthread_db, which deallocates its internal process handle and frees any other resources libthread_db has acquired with respect to the target process. <i>ta_p</i> specifies the target process to be deregistered.</p> <p>td_ta_get_ph() returns in *<i>ph_pp</i> the external process handle that corresponds to the internal process handle <i>ta_p</i>. This is useful for checking internal consistency.</p>												
RETURN VALUES	<table><tr><td>TD_OK</td><td>The call completed successfully.</td></tr><tr><td>TD_BADPH</td><td>A NULL external process handle was passed in to td_ta_new.</td></tr><tr><td>TD_ERR</td><td><i>ta_pp</i> is NULL, or an internal error occurred.</td></tr><tr><td>TD_DBERR</td><td>A call to one of the imported interface routines failed.</td></tr><tr><td>TD_MALLOC</td><td>Memory allocation failure.</td></tr><tr><td>TD_NOLIBTHREAD</td><td>The target process does not appear to be multithreaded.</td></tr></table>	TD_OK	The call completed successfully.	TD_BADPH	A NULL external process handle was passed in to td_ta_new.	TD_ERR	<i>ta_pp</i> is NULL, or an internal error occurred.	TD_DBERR	A call to one of the imported interface routines failed.	TD_MALLOC	Memory allocation failure.	TD_NOLIBTHREAD	The target process does not appear to be multithreaded.
TD_OK	The call completed successfully.												
TD_BADPH	A NULL external process handle was passed in to td_ta_new.												
TD_ERR	<i>ta_pp</i> is NULL, or an internal error occurred.												
TD_DBERR	A call to one of the imported interface routines failed.												
TD_MALLOC	Memory allocation failure.												
TD_NOLIBTHREAD	The target process does not appear to be multithreaded.												
ATTRIBUTES	See attributes(5) for description of the following attributes:												

td_ta_new(3THR)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO libthread_db(3THR), proc_service(3PROC), libthread_db(3LIB),
attributes(5)

td_ta_setconcurrency(3THR)

NAME | td_ta_setconcurrency – set concurrency level for target process

SYNOPSIS | cc [*flag ...*] *file ...* -lthread_db [*library ...*]

```
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_ta_setconcurrency(const td_thragent_t *ta_p, int
    level) ; ;
```

DESCRIPTION | td_ta_setconcurrency() sets the desired concurrency level for the process identified by *ta_p* to level, just as if a thread within the process had called thr_setconcurrency(). See thr_setconcurrency(3THR).

RETURN VALUES | TD_OK The call completed successfully.
 TD_BADTA An invalid internal process handle was passed in.
 TD_BADPH There is a NULL external process handle associated with this internal process handle. TD_NOCAPAB The client did not implement the ps_kill() routine in the imported interface. See ps_kill(3PROC).
 TD_DBERR A call to one of the imported interface routines failed.
 TD_ERR A libthread_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | libthread_db(3THR), ps_kill(3PROC), thr_setconcurrency(3THR), libthread_db(3LIB), attributes(5)

NAME	td_ta_sync_iter, td_ta_thr_iter, td_ta_tsd_iter – iterator functions on process handles from libthread_db
SYNOPSIS	<pre>cc [<i>flag ...</i>] <i>file ...</i> -lthread_db [<i>library ...</i>] #include <proc_service.h> #include <thread_db.h> typedef int td_sync_iter_f(const td_synchandle_t *sh_p, void *cbdata_p); typedef int td_thr_iter_f(const td_thrhandle_t *th_p, void *cbdata_p); typedef int td_key_iter_f(thread_key_t key, void (*destructor)(), void *cbdata_p); td_err_e td_ta_sync_iter(const td_thragent_t *ta_p, td_sync_iter_f *cb, void *cbdata_p); td_err_e td_ta_thr_iter(const td_thragent_t *ta_p, td_thr_iter_f *cb, void *cbdata_p, td_thr_state_e state, int ti_pri, sigset_t *ti_sigmask_p, unsigned ti_user_flags); td_err_e td_ta_tsd_iter(const td_thragent_t *ta_p, td_key_iter_f *cb, void *cbdata_p);</pre>
DESCRIPTION	<p>The <code>td_ta_sync_iter()</code>, <code>td_ta_thr_iter()</code>, and <code>td_ta_tsd_iter()</code> functions are iterator functions that when given a target process handle as an argument, return sets of handles for objects associated with the target process. The method is to call back a client-provided function once for each associated object, passing back a handle as well as the client-provided pointer <code>cb_data_p</code>. This enables a client to easily build a linked list of the associated objects. If the client-provided function returns non-zero, the iteration terminates, even if there are members remaining in the set of associated objects.</p> <p>The <code>td_ta_sync_iter()</code> function returns handles of synchronization objects (mutexes, readers-writer locks, semaphores, and condition variables) associated with a process. Some synchronization objects might not be known to <code>libthread_db</code> and will not be returned. If the process has initialized the synchronization object (by calling <code>mutex_init(3THR)</code>, for example) or a thread in the process has called a synchronization primitive (<code>mutex_lock()</code>, for example) using this object after <code>td_ta_new(3THR)</code> was called to attach to the process and <code>td_ta_sync_tracking_enable()</code> (see <code>td_sync_get_info(3THR)</code>) was called to enable synchronization object tracking, then a handle for the synchronization object will be passed to the callback function. See <code>td_sync_get_info(3THR)</code> for operations that can be performed on synchronization object handles.</p> <p>The <code>td_ta_thr_iter()</code> function returns handles for threads that are part of the target process. For <code>td_ta_thr_iter()</code>, the caller specifies several criteria to select a subset of threads for which the callback function should be called. Any of these</p>

td_ta_sync_iter(3THR)

selection criteria may be wild-carded. If all of them are wild-carded, then handles for all threads in the process will be returned.

The selection parameters and corresponding wild-card values are:

`state` (TD_THR_ANY_STATE):

Select only threads whose state matches `state`. See `td_thr_get_info(3THR)` for a list of thread states.

`ti_pri` (TD_THR_LOWEST_PRIORITY):

Select only threads for which the priority is at least `ti_pri`.

`ti_sigmask_p` (TD_SIGNO_MASK):

Select only threads whose signal mask exactly matches `*ti_sigmask_p`.

`ti_user_flags` (TD_THR_ANY_USER_FLAGS):

Select only threads whose user flags (specified at thread creation time) exactly match `ti_user_flags`.

The `td_ta_tsd_iter()` function returns the thread-specific data keys in use by the current process. Thread-specific data for a particular thread and key can be obtained by calling `td_thr_tsd(3THR)`.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTA	An invalid process handle was passed in.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	The call did not complete successfully.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO

`libthread_db(3LIB)`, `libthread_db(3THR)`, `mutex_init(3THR)`, `td_sync_get_info(3THR)`, `td_thr_get_info(3THR)`, `td_thr_tsd(3THR)`, `attributes(5)`

NAME td_thr_dbsuspend, td_thr_dbresume – suspend and resume threads in libthread_db

SYNOPSIS

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

#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_dbsuspend(const td_thrhandle_t *th_p);
td_err_e td_thr_dbresume(const td_thrhandle_t *th_p);
```

DESCRIPTION

These operations suspend and resume the thread identified by *th_p*. A thread that has been suspended with `td_thr_dbsuspend()` is said to be in the "dbsuspended" state. A thread whose "dbsuspended" flag is set will not execute. If an unbound thread enters the "dbsuspended" state and is currently assigned to a lightweight process (LWP), then the LWP becomes available for assignment to a different thread.

A thread's "dbsuspended" state is independent of the suspension state controlled by calls to `thr_suspend(3THR)` and `thr_continue(3THR)` from within the target process. Calling `thr_continue(3THR)` within the target process on a thread that has been suspended during a call to `td_thr_dbsuspend()` will not cause that thread to resume execution; only a call to `td_thr_dbresume()` will do that.

RETURN VALUES

TD_OK The call completed successfully.

TD_BADTH An invalid thread handle was passed in.

TD_DBERR A call to one of the imported interface routines failed.

TD_NOCAPAB The "agent thread" in the target process has not completed initialization, so this operation cannot be performed. The operation can be performed after the target process has been allowed to make some forward progress. See also `libthread_db(3THR)`

TD_ERR A libthread_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libthread_db(3THR)`, `thr_continue(3THR)`, `thr_suspend(3THR)`, `libthread_db(3LIB)`, `attributes(5)`

td_thr_getgregs(3THR)

NAME	td_thr_getgregs, td_thr_setgregs, td_thr_getfpregs, td_thr_setfpregs, td_thr_getxregsize, td_thr_getxregs, td_thr_setxregs – reading and writing thread registers in libthread_db
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i> ... -lthread_db [<i>library</i> ...] #include <proc_service.h> #include <thread_db.h> td_err_e td_thr_getgregs(const td_thrhandle_t *th_p, pgregset_t tgregset); td_err_e td_thr_setgregs(const td_thrhandle_t *th_p, pgregset_t tgregset); td_err_e td_thr_getfpregs(const td_thrhandle_t *th_p, prfpregset_t *fpregset); td_err_e td_thr_setfpregs(const td_thrhandle_t *th_p, prfpregset_t *fpregset); td_err_e td_thr_getxregsize(const td_thrhandle_t *th_p, int *xregsize); td_err_e td_thr_getxregs(const td_thrhandle_t *th_p, prxregset_t *xregset); td_err_e td_thr_setxregs(const td_thrhandle_t *th_p, prxregset_t *xregset);</pre>
DESCRIPTION	<p>These routines read and write the register sets associated with thread <i>th_p</i>. <code>td_thr_getgregs()</code> and <code>td_thr_setgregs()</code> get and set, respectively, the general registers of thread <i>th_p</i>. <code>td_thr_getfpregs()</code> and <code>td_thr_setfpregs()</code> get and set, respectively, the thread's floating point register set. <code>td_thr_getxregsize()</code>, <code>td_thr_getxregs()</code>, and <code>td_thr_setxregs()</code> are SPARC-specific. <code>td_thr_getxregsize()</code> returns in <i>*xregsize</i> the size of the architecture-dependent extra state registers. <code>td_thr_getxregs()</code> and <code>td_thr_setxregs()</code> get and set, respectively, those extra state registers. On non-SPARC architectures, these routines return <code>TD_NOXREGS</code>.</p> <p>If thread <i>th_p</i> is currently executing on a lightweight process (LWP), these routines will read or write, respectively, the appropriate register set to the LWP using the imported interface. If the thread is not currently executing on a LWP, then the floating point and extra state registers may not be read or written. Some of the general registers may also not be readable or writable, depending on the architecture. In this case, <code>td_thr_getfpregs()</code> and <code>td_thr_setfpregs()</code> will return <code>TD_NOFPREGS</code>, and <code>td_thr_getxregs()</code> and <code>td_thr_setxregs()</code> will return <code>TD_NOXREGS</code>. Calls to <code>td_thr_getgregs()</code> and <code>td_thr_setgregs()</code> will succeed, but values returned for unreadable registers will be undefined, and values specified for unwritable registers will be ignored. In this instance, a value of</p>

td_thr_getgregs(3THR)

TD_PARTIALREGS will be returned. See the architecture-specific notes that follow regarding the registers that may be read and written for a thread not currently executing on a LWP.

SPARC On a thread not currently assigned to a LWP, only %i0-%i7, %l0-%l7, %g7, %pc, and %sp (%o6) may be read or written. %pc and %sp refer to the program counter and stack pointer that the thread will have when it resumes execution.

Intel IA On a thread not currently assigned to a LWP, only %pc, %sp, %ebp, %edi, %edi, and %ebx may be read.

RETURN VALUES

- TD_OK The call completed successfully.
- TD_BADTH An invalid thread handle was passed in.
- TD_DBERR A call to one of the imported interface routines failed.
- TD_PARTIALREGS Because the thread is not currently assigned to a LWP, not all registers were read or written. See DESCRIPTION for a discussion about which registers are not saved when a thread is not assigned to an LWP.
- TD_NOFPREGS Floating point registers could not be read or written, either because the thread is not currently assigned to an LWP, or because the architecture does not have such registers.
- TD_NOXREGS Architecture-dependent extra state registers could not be read or written, either because the thread is not currently assigned to an LWP, or because the architecture does not have such registers, or because the architecture is not a SPARC architecture.
- TD_ERR A libthread_db internal error occurred.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO

libthread_db(3THR), libthread_db(3LIB), attributes(5)

td_thr_get_info(3THR)

NAME	td_thr_get_info – get thread information in libthread_db library of interfaces
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i> ... -lthread_db [<i>library</i> ...] #include <proc_service.h> #include <thread_db.h> td_err_e td_thr_get_info(const td_thrhandle_t *<i>th_p</i>, td_thrinfo_t *<i>ti_p</i>);</pre>
DESCRIPTION	<p>The <code>td_thr_get_info()</code> function fills in the <code>td_thrinfo_t</code> structure <code>*ti_p</code> with values for the thread identified by <code>th_p</code>.</p> <p>The <code>td_thrinfo_t</code> structure contains the following fields:</p> <pre>typedef struct td_thrinfo_t { td_thragen_tx *ti_ta_p /* internal process handle */ unsigned ti_user_flags; /* value of flags parameter */ thread_t ti_tid; /* thread identifier */ char *ti_tls; /* pointer to thread-local storage*/ paddr ti_startfunc; /* address of function at which thread execution began*/ paddr ti_stkbase; /* base of thread's stack area*/ int ti_stksize; /* size in bytes of thread's allocated stack region*/ paddr ti_ro_area; /* address of utthread_t structure*/ int ti_ro_size /* size of the utthread_t structure in bytes */ td_thr_state_e ti_state /* state of the thread */ uchar_t ti_db_suspended /* non-zero if thread suspended by td_thr_dbsuspend*/ td_thr_type_e ti_type /* type of the thread*/ int ti_pc /* value of thread's program counter*/ int ti_sp /* value of thread's stack counter*/ short ti_flags /* set of special flags used by libthread*/ int ti_pri /* priority of thread returned by thr_getprio(3T)*/ lwpid_t ti_lid /* id of light weight process (LWP) executing this thread*/ sigset_t ti_sigmask /* thread's signal mask. See thr_sigsetmask(3T)*/ u_char ti_traceme /* non-zero if event tracing is on*/ u_char_t ti_preemptflag /* non-zero if thread preempted when last active*/ u_char_t ti_pirecflag /* non-zero if thread runs priority beside regular */ sigset_t ti_pending /* set of signals pending for this thread*/ td_thr_events_t ti_events /* bitmap of events enabled for this thread*/ };</pre> <p><code>td_thragent_t *ti_ta_p</code> is the internal process handle identifying the process of which the thread is a member.</p>

td_thr_get_info(3THR)

unsigned *ti_user_flags* is the value of the *flags* parameter passed to `thr_create(3THR)` when the thread was created.

`thread_t` *ti_tid* is the thread identifier for the thread returned by `libthread` when created with `thr_create(3THR)`.

`char *ti_tls` is the thread's pointer to thread-local storage.

`psaddr_t` *ti_startfunc* is the address of the function at which thread execution began, as specified when the thread was created with `thr_create(3THR)`.

`psaddr_t` *ti_stkbase* is the base of the thread's stack area.

`int` *ti_stksize* is the size in bytes of the thread's allocated stack region.

`psaddr_t` *ti_ro_area* is the address of the `libthread`-internal `uthread_t` structure for this thread. Since accessing the `uthread_t` structure directly violates the encapsulation provided by `libthread_db`, this field should generally not be used. However, it may be useful as a prototype for extensions.

`td_thr_state_e` *ti_state* is the state in which the thread is. The `td_thr_state_e` enumeration type may contain the following values:

<code>TD_THR_ANY_STATE</code>	Never returned by <code>td_thr_get_info</code> . <code>TD_THR_ANY_STATE</code> is used as a wildcard to select threads in <code>td_ta_thr_iter()</code> .
<code>TD_THR_UNKNOWN</code>	<code>libthread_db</code> cannot determine the state of the thread.
<code>TD_THR_STOPPED</code>	The thread has been stopped by a call to <code>thr_suspend(3THR)</code> .
<code>TD_THR_RUN</code>	The thread is runnable, but it is not currently assigned to a LWP.
<code>TD_THR_ACTIVE</code>	The thread is currently executing on a LWP.
<code>TD_THR_ZOMBIE</code>	The thread has exited, but it has not yet been deallocated by a call to <code>thr_join(3THR)</code> .
<code>TD_THR_SLEEP</code>	The thread is not currently runnable.
<code>TD_THR_STOPPED_ASLEEP</code>	The thread is both blocked by <code>TD_THR_SLEEP</code> , and stopped by a call to <code>td_thr_dbsuspend(3THR)</code> .

`uchar_t` *ti_db_suspended* is non-zero if and only if this thread is currently suspended because the controlling process has called `td_thr_dbsuspend` on it.

td_thr_get_info(3THR)

`td_thr_type_e ti_type` is a type of thread. It will be either `TD_THR_USER` for a user thread (one created by the application), or `TD_THR_SYSTEM` for one created by `libthread`.

`int ti_pc` is the value of the thread's program counter, provided that the thread's `ti_state` value is `TD_THR_SLEEP`, `TD_THR_STOPPED`, or `TD_THR_STOPPED_ASLEEP`. Otherwise, the value of this field is undefined.

`int ti_sp` is the value of the thread's stack pointer, provided that the thread's `ti_state` value is `TD_THR_SLEEP`, `TD_THR_STOPPED`, or `TD_THR_STOPPED_ASLEEP`. Otherwise, the value of this field is undefined.

`short ti_flags` is a set of special flags used by `libthread`, currently of use only to those debugging `libthread`.

`int ti_pri` is the thread's priority, as it would be returned by `thr_getprio(3THR)`.

`lwpid_t ti_lid` is the ID of the LWP executing this thread, or the ID of the LWP that last executed this thread, if this thread is not currently assigned to a LWP.

`sigset_t ti_sigmask` is this thread's signal mask. See `thr_sigsetmask(3THR)`.

`u_char ti_traceme` is non-zero if and only if event tracing for this thread is on.

`uchar_t ti_preemptflag` is non-zero if and only if the thread was preempted the last time it was active.

`uchar_t ti_pirecflag` is non-zero if and only if due to priority inheritance the thread is currently running at a priority other than its regular priority.

`td_thr_events_t ti_events` is the bitmap of events enabled for this thread.

RETURN VALUES

<code>TD_OK</code>	The call completed successfully.
<code>TD_BADTH</code>	An invalid thread handle was passed in.
<code>TD_DBERR</code>	A call to one of the imported interface routines failed.
<code>TD_ERR</code>	The call did not complete successfully.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO

`libthread(3LIB)`, `libthread_db(3LIB)`, `libthread_db(3THR)`, `td_ta_thr_iter(3THR)`, `td_thr_dbsuspend(3THR)`, `thr_create(3THR)`, `thr_getprio(3THR)`, `thr_join(3THR)`, `thr_sigsetmask(3THR)`, `thr_suspend(3THR)`, `threads(3THR)`, `attributes(5)`

NAME	td_thr_lockowner – iterate over the set of locks owned by a thread										
SYNOPSIS	<pre>cc [<i>flag ...</i>] <i>file ...</i> -lthread_db [<i>library ...</i>] #include <proc_service.h> #include <thread_db.h> td_err_e td_thr_lockowner(const td_thrhandle_t *<i>th_p</i>, td_sync_iter_f *<i>cb</i>, void *<i>cb_data_p</i>);</pre>										
DESCRIPTION	<p>td_thr_lockowner() calls the iterator function <i>cb</i> once for every mutex that is held by the thread whose handle is <i>th_p</i>. The synchronization handle and the pointer <i>cb_data_p</i> are passed to the function. See td_ta_thr_iter(3THR) for a similarly structured function.</p> <p>Iteration terminates early if the callback function <i>cb</i> returns a non-zero value.</p>										
RETURN VALUES	<table border="0"> <tr> <td style="padding-right: 20px;">TD_OK</td> <td>The call completed successfully.</td> </tr> <tr> <td>TD_BADTH</td> <td>An invalid thread handle was passed in.</td> </tr> <tr> <td>TD_BADPH</td> <td>There is a NULL external process handle associated with this internal process handle.</td> </tr> <tr> <td>TD_DBERR</td> <td>A call to one of the imported interface routines failed.</td> </tr> <tr> <td>TD_ERR</td> <td>A libthread_db internal error occurred.</td> </tr> </table>	TD_OK	The call completed successfully.	TD_BADTH	An invalid thread handle was passed in.	TD_BADPH	There is a NULL external process handle associated with this internal process handle.	TD_DBERR	A call to one of the imported interface routines failed.	TD_ERR	A libthread_db internal error occurred.
TD_OK	The call completed successfully.										
TD_BADTH	An invalid thread handle was passed in.										
TD_BADPH	There is a NULL external process handle associated with this internal process handle.										
TD_DBERR	A call to one of the imported interface routines failed.										
TD_ERR	A libthread_db internal error occurred.										
ATTRIBUTES	See attributes(5) for description of the following attributes:										
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe						
ATTRIBUTE TYPE	ATTRIBUTE VALUE										
MT-Level	Safe										
SEE ALSO	libthread_db(3THR), td_ta_thr_iter(3THR), libthread_db(3LIB), attributes(5)										

td_thr_setprio(3THR)

NAME | td_thr_setprio – set the priority of a thread

SYNOPSIS | cc [*flag ...*] *file ...* -lthread_db [*library ...*]

```
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_setprio(const td_thrhandle_t *th_p, const int
    new_prio);
```

DESCRIPTION | td_thr_setprio() sets thread *th_p*'s priority to *new_prio*, just as if a thread within the process had called thr_setprio(). See thr_setprio(3THR).

RETURN VALUES | TD_OK The call completed successfully.
TD_BADTH An invalid thread handle was passed in.
TD_DBERR A call to one of the imported interface routines failed.
TD_ERR *new_prio* is an illegal value (out of range).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | libthread_db(3THR), thr_setprio(3THR), libthread_db(3LIB), attributes(5)

td_thr_setsigpending(3THR)

NAME td_thr_setsigpending, td_thr_sigsetmask – manage thread signals for libthread_db

SYNOPSIS

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

#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_setsigpending(const td_thrhandle_t * th_p, const
    uchar_t ti_sigpending_flag, const sigset_t ti_sigmask);

td_err_e td_thr_sigsetmask(const td_thrhandle_t *th_p, const
    sigset_t ti_sigmask);
```

DESCRIPTION

The `td_thr_setsigpending()` and `td_thr_setsigmask()` operations affect the signal state of the thread identified by *th_p*.

`td_thr_setsigpending()` sets the set of pending signals for thread *th_p* to *ti_sigpending*. The value of the libthread-internal field that indicates whether a thread has any signal pending is set to *ti_sigpending_flag*. To be consistent, *ti_sigpending_flag* should be zero if and only if all of the bits in *ti_sigpending* are zero.

`td_thr_sigsetmask()` sets the signal mask of the thread *th_p* as if the thread had set its own signal mask by way of `thr_sigsetmask(3THR)`. The new signal mask is the value of *ti_sigmask*.

There is no equivalent to the `SIG_BLOCK` or `SIG_UNBLOCK` operations of `thr_sigsetmask(3THR)`, which mask or unmask specific signals without affecting the mask state of other signals. To block or unblock specific signals, either stop the whole process, or the thread, if necessary, by `td_thr_dbsuspend()`. Then determine the thread's existing signal mask by calling `td_thr_get_info()` and reading the *ti_sigmask* field of the `td_thrinfo_t` structure returned. Modify it as desired, and set the new signal mask with `td_thr_sigsetmask()`.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	A libthread_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libthread_db(3THR)`, `td_thr_dbsuspend(3THR)`, `td_thr_get_info(3THR)`, `libthread_db(3LIB)`, `attributes(5)`

td_thr_sleepinfo(3THR)

NAME td_thr_sleepinfo – return the synchronization handle for the object on which a thread is blocked

SYNOPSIS

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

#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_sleepinfo(const td_thrhandle_t *th_p,
    td_synchandle_t *sh_p);
```

DESCRIPTION td_thr_sleepinfo() returns in *sh_p the handle of the synchronization object on which a sleeping thread is blocked.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	The thread th_p is not blocked on a synchronization object, or a libthread_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO libthread_db(3THR), libthread_db(3LIB), attributes(5)

td_thr_tsd(3THR)

NAME | td_thr_tsd – get a thread’s thread-specific data for libthread_db library of interfaces

SYNOPSIS |

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

|

```
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_tsd(const td_thrhandle_t, const thread_key_t key,
void *data_pp);
```

DESCRIPTION | td_thr_tsd() returns in **data_pp* the thread-specific data pointer for the thread identified by *th_p* and the thread-specific data key *key*. This is the same value that thread *th_p* would obtain if it called thr_getspecific(3THR).

| To find all the thread-specific data keys in use in a given target process, call td_ta_tsd_iter(3THR).

RETURN VALUES | TD_OK The call completed successfully.

| TD_BADTH An invalid thread handle was passed in.

| TD_DBERR A call to one of the imported interface routines failed.

| TD_ERR A libthread_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | libthread_db(3THR), td_ta_tsd_iter(3THR), thr_getspecific(3THR), libthread_db(3LIB), attributes(5)

td_thr_validate(3THR)

NAME | td_thr_validate – test a thread handle for validity

SYNOPSIS |

```
cc [ flag ... ] file ... -lthread_db [ library ... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_validate(const td_thrhandle_t *th_p);
```

DESCRIPTION | td_thr_validate() tests whether *th_p* is a valid thread handle. A valid thread handle may become invalid if its thread exits.

RETURN VALUES | TD_OK The call completed successfully. *th_p* is a valid thread handle.
TD_BADTH *th_p* was NULL.
TD_DBERR A call to one of the imported interface routines failed.
TD_NOTHR *th_p* is not a valid thread handle.
TD_ERR A libthread_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | libthread_db(3THR), libthread_db(3LIB), attributes(5)

NAME	thr_create – create a thread		
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i>[<i>library...</i>] #include <thread.h> int thr_create(void *<i>stack_base</i>, size_t <i>stack_size</i>, void *(*<i>start_func</i>) (void*), void *<i>arg</i>, long <i>flags</i>, thread_t *<i>new_thread_ID</i>);</pre>		
DESCRIPTION	<p>Thread creation adds a new thread of control to the current process. The procedure <code>main()</code> is a single thread of control. Each thread executes concurrently with all other threads within the calling process and with other threads from other active processes.</p> <p>Although a newly created thread shares all of the calling process's global data with the other threads in the process, it has its own set of attributes and private execution stack. The new thread inherits the calling thread's signal mask, possibly, and scheduling priority. Pending signals for a new thread are not inherited and will be empty.</p> <p>The call to create a thread takes the address of a user-defined function, specified by <i>start_func</i>, as one of its arguments. This function is the complete execution routine for the new thread.</p> <p>The lifetime of a thread begins with the successful return from <code>thr_create()</code>, which calls <i>start_func()</i> and ends with one of the following:</p> <ul style="list-style-type: none"> ■ the normal completion of <i>start_func()</i>, ■ the return from an explicit call to <code>thr_exit(3THR)</code>, or ■ the conclusion of the calling process (see <code>exit(2)</code>). <p>The new thread performs by calling the function defined by <i>start_func</i> with only one argument, <i>arg</i>. If more than one argument needs to be passed to <i>start_func</i>, the arguments can be packed into a structure, the address of which can be passed to <i>arg</i>.</p> <p>If <i>start_func</i> returns, the thread terminates with the exit status set to the <i>start_func</i> return value (see <code>thr_exit(3THR)</code>).</p> <p>When the thread from which <code>main()</code> originated returns, the effect is the same as if an implicit call to <code>exit()</code> were made using the return value of <code>main()</code> as the exit status. This behavior differs from a <i>start_func</i> return. If <code>main()</code> calls <code>thr_exit(3THR)</code>, only the <code>main</code> thread exits, not the entire process.</p> <p>If the thread creation fails, a new thread is not created and the contents of the location referenced by the pointer to the new thread are undefined.</p> <p>The <i>flags</i> argument specifies which attributes are modifiable for the created thread. The value in <i>flags</i> is determined by the bitwise inclusive-OR of the following:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">THR_BOUND</td> <td>This flag affects the <code>contentionscope</code> attribute of the thread. The new thread is created permanently bound to an LWP (that is, it is a <i>bound thread</i>.)</td> </tr> </table>	THR_BOUND	This flag affects the <code>contentionscope</code> attribute of the thread. The new thread is created permanently bound to an LWP (that is, it is a <i>bound thread</i> .)
THR_BOUND	This flag affects the <code>contentionscope</code> attribute of the thread. The new thread is created permanently bound to an LWP (that is, it is a <i>bound thread</i> .)		

thr_create(3THR)

THR_DETACHED	This flag affects the detachstate attribute of the thread. The new thread is created detached. The exit status of a detached thread is not accessible to other threads. Its thread ID and other resources may be re-used as soon as the thread terminates. <code>thr_join(3THR)</code> will not wait for a detached thread.
THR_NEW_LWP	This flag is obsolete and is maintained for compatibility.
THR_SUSPENDED	This flag affects the suspended attribute of the thread. The new thread is created suspended and will not execute <code>start_func</code> until it is started by <code>thr_continue()</code> .
THR_DAEMON	This flag affects the daemon attribute of the thread. In addition to being created detached (THR_DAEMON implies THR_DETACHED), the thread is marked as a daemon. Daemon threads do not interfere with the exit conditions for a process. A process will terminate when the last non-daemon thread exits or the process calls <code>exit(2)</code> . Also, a thread that is waiting in <code>thr_join()</code> for any thread to terminate will return EDEADLK when all remaining threads in the process are either daemon threads or other threads waiting in <code>thr_join()</code> . Daemon threads are most useful in libraries that want to use threads.

Default thread creation:

```
thread_t tid;
void *start_func(void *), *arg;
thr_create(NULL, NULL, start_func, arg, NULL, &tid);
```

User-defined thread creation (create a bound permanently to an LWP, that is, a bound thread):

```
thr_create(NULL, NULL, start_func, arg, THR_BOUND, &tid);
```

With `thr_create()`, the new thread uses the stack beginning at the address specified by `stack_base` and continuing for `stack_size` bytes. The `stack_size` argument must be greater than the value returned by `thr_min_stack(3THR)`. If `stack_base` is NULL, `thr_create()` allocates a stack for the new thread with at least `stack_size` bytes. If `stack_size` is 0, a default size is used. If `stack_size` is not 0, it must be greater than the value returned by `thr_min_stack(3THR)`. See NOTES.

When `new_thread_ID` is not NULL, it points to a location where the ID of the new thread is stored if `thr_create()` is successful. The ID is only valid within the calling process.

RETURN VALUES

If successful, the `thr_create()` function returns 0. Otherwise, an error value is returned to indicate the error. If the application is not linked with the threads library, -1 is returned.

ERRORS	EAGAIN	The system-imposed limit on the total number of threads in a process has been exceeded or some system resource has been exceeded (for example, too many LWPs were created).
	EINVAL	The <i>stack_base</i> argument is not NULL and <i>stack_size</i> is less than the value returned by <code>thr_min_stack(3THR)</code> , or the <i>stack_base</i> argument is NULL and <i>stack_size</i> is not 0 and is less than the value returned by <code>thr_min_stack(3THR)</code> .
	ENOMEM	The system cannot allocate stack for the thread.

The `thr_create()` function may use `mmap()` to allocate thread stacks from `MAP_PRIVATE`, `MAP_NORESERVE`, and `MAP_ANON` memory mappings if *stack_base* is NULL, and consequently may return upon failure the relevant error values returned by `mmap()`. See the `mmap(2)` manual page for these error values.

EXAMPLES **EXAMPLE 1** This is an example of concurrency with multi-threading. Since POSIX threads and Solaris threads are fully compatible even within the same process, this example uses `pthread_create()` if you execute `a.out 0`, or `thr_create()` if you execute `a.out 1`.

Five threads are created that simultaneously perform a time-consuming function, `sleep(10)`. If the execution of this process is timed, the results will show that all five individual calls to `sleep` for ten-seconds completed in about ten seconds, even on a uniprocessor. If a single-threaded process calls `sleep(10)` five times, the execution time will be about 50-seconds.

The command-line to time this process is:

```
/usr/bin/time a.out 0 (for POSIX threading)
```

or

```
/usr/bin/time a.out 1 (for Solaris threading)
```

```
/* cc thisfile.c -lthread -lpthread */
#define _REENTRANT /* basic 3-lines for threads */
#include <pthread.h>
#include <thread.h>
#define NUM_THREADS 5
#define SLEEP_TIME 10

void *sleeping(void *); /* thread routine */
int i;
thread_t tid[NUM_THREADS]; /* array of thread IDs */

int
main(int argc, char *argv[])
{
    if (argc == 1) {
        printf("use 0 as arg1 to use pthread_create( )\n");
        printf("or use 1 as arg1 to use thr_create( )\n");
        return (1);
    }
}
```

thr_create(3THR)

EXAMPLE 1 This is an example of concurrency with multi-threading. Since POSIX threads and Solaris threads are fully compatible even within the same process, this example uses `pthread_create()` if you execute `a.out 0`, or `thr_create()` if you execute `a.out 1`.
(Continued)

```
switch (*argv[1]) {
case '0': /* POSIX */
    for ( i = 0; i < NUM_THREADS; i++)
        pthread_create(&tid[i], NULL, sleeping,
            (void *)SLEEP_TIME);
    for ( i = 0; i < NUM_THREADS; i++)
        pthread_join(tid[i], NULL);
    break;

case '1': /* Solaris */
    for ( i = 0; i < NUM_THREADS; i++)
        thr_create(NULL, 0, sleeping, (void *)SLEEP_TIME, 0,
            &tid[i]);
    while (thr_join(0, NULL, NULL) == 0)
        ;
    break;
} /* switch */
printf("main( ) reporting that all %d threads have terminated\n", i);
return (0);
} /* main */

void *
sleeping(void *arg)
{
    int sleep_time = (int)arg;
    printf("thread %d sleeping %d seconds ...\n", thr_self( ), sleep_time);
    sleep(sleep_time);
    printf("\nthread %d awakening\n", thr_self( ));
    return (NULL);
}
```

Had `main()` not waited for the completion of the other threads (using `pthread_join(3THR)` or `thr_join(3THR)`), it would have continued to process concurrently until it reached the end of its routine and the entire process would have exited prematurely (see `exit(2)`).

EXAMPLE 2 Creating a default thread with a new signal mask.

The following example demonstrates how to create a default thread with a new signal mask. The *new_mask* argument is assumed to have a value different from the creator's signal mask (*orig_mask*). The *new_mask* argument is set to block all signals except for `SIGINT`. The creator's signal mask is changed so that the new thread inherits a different mask, and is restored to its original value after `thr_create()` returns.

This example assumes that `SIGINT` is also unmasked in the creator. If it is masked by the creator, then unmasking the signal opens the creator to this signal. The other alternative is to have the new thread set its own signal mask in its start routine.

EXAMPLE 2 Creating a default thread with a new signal mask. (Continued)

```

thread_t tid;
sigset_t new_mask, orig_mask;
int error;

(void) sigfillset(&new_mask);
(void) sigdelset(&new_mask, SIGINT);
(void) thr_sigsetmask(SIG_SETMASK, &new_mask, &orig_mask);
error = thr_create(NULL, 0, do_func, NULL, 0, &tid);
(void) thr_sigsetmask(SIG_SETMASK, &orig_mask, NULL);

```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO _lwp_create(2), exit(2), getrlimit(2), mmap(2), exit(3C), sleep(3C), thr_exit(3THR), thr_join(3THR), thr_min_stack(3THR), thr_setconcurrency(3THR), thr_suspend(3THR), threads(3THR), attributes(5), standards(5)

NOTES Since multithreaded-application threads execute independently of each other, their relative behavior is unpredictable. It is therefore possible for the thread executing main() to finish before all other user-application threads.

Using thr_join(3THR) in the following syntax,

```

while (thr_join(0, NULL, NULL) == 0);

```

will cause the invoking thread (which may be main()) to wait for the termination of all non-daemon threads, excluding threads that are themselves waiting in thr_join(); however, the second and third arguments to thr_join() need not necessarily be NULL.

A thread has not terminated until thr_exit() has finished. The only way to determine this is by thr_join(). When thr_join() returns a departed thread, it means that this thread has terminated and its resources are reclaimable. For instance, if a user specified a stack to thr_create(), this stack can only be reclaimed after thr_join() has reported this thread as a departed thread. It is not possible to determine when a *detached* thread has terminated. A detached thread disappears without leaving a trace.

Typically, thread stacks allocated by thr_create() begin on page boundaries and any specified (a red-zone) size is rounded up to the next page boundary. A page with no access permission is appended to the top of the stack so that most stack overflows will result in a SIGSEGV signal being sent to the offending thread. Thread stacks allocated by the caller are used as is.

thr_create(3THR)

Using a default stack size for the new thread, instead of passing a user-specified stack size, results in much better `thr_create()` performance. The default stack size for a user-thread is 1 megabyte in a 32-bit process and 2 megabyte in a 64-bit process.

A user-specified stack size must be greater than the value `THR_MIN_STACK`. A minimum stack size may not accommodate the stack frame for the user thread function `start_func`. If a stack size is specified, it must accommodate `start_func` requirements and the functions that it may call in turn, in addition to the minimum requirement.

It is usually very difficult to determine the runtime stack requirements for a thread. `THR_MIN_STACK` specifies how much stack storage is required to execute a NULL `start_func`. The total runtime requirements for stack storage are dependent on the storage required to do runtime linking, the amount of storage required by library runtimes (like `printf()`) that your thread calls. Since these storage parameters are not known before the program runs, it is best to use default stacks. If you know your runtime requirements or decide to use stacks that are larger than the default, then it makes sense to specify your own stacks.

NAME	threads, pthreads – concepts related to POSIX pthreads and Solaris threads and the libpthread and libthread libraries				
POSIX	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h></pre>				
Solaris	<pre>cc - mt [flag...] file...[library...] #include <sched.h> #include <thread.h></pre>				
DESCRIPTION	<p>POSIX and Solaris threads each have their own implementation of the threads library. The libpthread library is associated with POSIX; the libthread library is associated with Solaris. Both implementations are interoperable, their functionality similar, and can be used within the same application. Only POSIX threads are guaranteed to be fully portable to other POSIX-compliant environments. POSIX and Solaris threads require different source, include files and linking libraries. See SYNOPSIS.</p>				
Similarities	<p>Most of the functions in the libpthread and libthread, libraries have a counterpart in the other corresponding library. POSIX function names, with the exception of the semaphore names, have a "pthread" prefix. Function names for similar POSIX and Solaris have similar endings. Typically, similar POSIX and Solaris functions have the same number and use of arguments.</p>				
Differences	<p>POSIX pthreads and Solaris threads differ in the following ways:</p> <ul style="list-style-type: none"> ■ POSIX threads are more portable. ■ POSIX threads establish characteristics for each thread according to configurable attribute objects. ■ POSIX pthreads implement thread cancellation. ■ POSIX pthreads enforce scheduling algorithms. ■ POSIX pthreads allow for clean-up handlers for fork(2) calls. ■ Solaris threads can be suspended and continued. ■ Solaris threads implement an optimized mutex and interprocess robust mutex locks. ■ Solaris threads implement daemon threads, for whose demise the process does not wait. 				
Function Comparison	<p>The following table compares the POSIX pthreads and Solaris threads functions. When a comparable interface is not available either in POSIX pthreads or Solaris threads, a hyphen (-) appears in the column.</p>				
Functions Related to Creation	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 50%;"></th> <th style="width: 50%; text-align: center;">Solaris (libthread)</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">POSIX (libpthread)</td> <td></td> </tr> </tbody> </table>		Solaris (libthread)	POSIX (libpthread)	
	Solaris (libthread)				
POSIX (libpthread)					

threads(3THR)

	pthread_create()	thr_create()
	pthread_attr_init()	-
	pthread_attr_setdetachstate()	-
	pthread_attr_getdetachstate()	-
	pthread_attr_setinheritsched()	-
	pthread_attr_getinheritsched()	-
	pthread_attr_setschedparam()	-
	pthread_attr_getschedparam()	-
	pthread_attr_setschedpolicy()	-
	pthread_attr_getschedpolicy()	-
	pthread_attr_setscope()	-
	pthread_attr_getscope()	-
	pthread_attr_setstackaddr()	-
	pthread_attr_getstackaddr()	-
	pthread_attr_setstacksize()	-
	pthread_attr_getstacksize()	-
	pthread_attr_getguardsize()	-
	pthread_attr_setguardsize()	-
	pthread_attr_destroy()	-
	-	thr_min_stack()
	<hr/>	
Functions Related to Exit	POSIX (libpthread)	Solaris (libthread)
	pthread_exit()	thr_exit()
	pthread_join()	thr_join()
	pthread_detach()	-
	<hr/>	
Functions Related to Thread Specific Data	POSIX (libpthread)	Solaris (libthread)
	pthread_key_create()	thr_keycreate()
	pthread_setspecific()	thr_setspecific()
	pthread_getspecific()	thr_getspecific()

	pthread_key_delete()	-
Functions Related to Signals	POSIX (libpthread)	Solaris (libthread)
	pthread_sigmask()	thr_sigsetmask()
	pthread_kill()	thr_kill()
Functions Related to IDs	POSIX (libpthread)	Solaris (libthread)
	pthread_self()	thr_self()
	pthread_equal()	-
	-	thr_main()
Functions Related to Scheduling	POSIX (libpthread)	Solaris (libthread)
	-	thr_yield()
	-	thr_suspend()
	-	thr_continue()
	pthread_setconcurrency()	thr_setconcurrency()
	pthread_getconcurrency()	thr_getconcurrency()
	pthread_setschedparam()	thr_setprio()
	pthread_getschedparam()	thr_getprio()
Functions Related to Cancellation	POSIX (libpthread)	Solaris (libthread)
	pthread_cancel()	-
	pthread_setcancelstate()	-
	pthread_setcanceltype()	-
	pthread_testcancel()	-
	pthread_cleanup_pop()	-
	pthread_cleanup_push()	-
Functions Related to Mutexes	POSIX (libpthread)	Solaris (libthread)
	pthread_mutex_init()	mutex_init()
	pthread_mutexattr_init()	-

threads(3THR)

pthread_mutexattr_setpshared()	-
pthread_mutexattr_getpshared()	-
pthread_mutexattr_setprotocol()	-
pthread_mutexattr_getprotocol()	-
pthread_mutexattr_setprioceiling()	-
pthread_mutexattr_getprioceiling()	-
pthread_mutexattr_settype()	-
pthread_mutexattr_gettype()	-
pthread_mutexattr_destroy()	-
pthread_mutex_setprioceiling()	-
pthread_mutex_getprioceiling()	-
pthread_mutex_lock()	mutex_lock()
pthread_mutex_trylock()	mutex_trylock()
pthread_mutex_unlock()	mutex_unlock()
pthread_mutex_destroy()	mutex_destroy()

Functions Related to Condition Variables

	POSIX (libpthread)	Solaris (libthread)
pthread_cond_init()		cond_init()
pthread_condattr_init()		-
pthread_condattr_setpshared()		-
pthread_condattr_getpshared()		-
pthread_condattr_destroy()		-
pthread_cond_wait()		cond_wait()
pthread_cond_timedwait()		cond_timedwait()
pthread_cond_signal()		cond_signal()
pthread_cond_broadcast()		cond_broadcast()
pthread_cond_destroy()		cond_destroy()

Functions Related to Reader/Writer Locking

	POSIX (libpthread)	Solaris (libthread)
pthread_rwlock_init()		rwlock_init()

	pthread_rwlock_rdlock()	rw_rdlock()
	pthread_rwlock_tryrdlock()	rw_tryrdlock()
	pthread_rwlock_wrlock()	rw_wrlock()
	pthread_rwlock_trywrlock()	rw_trywrlock()
	pthread_rwlock_unlock()	rw_unlock()
	pthread_rwlock_destroy()	rwlock_destroy()
	pthread_rwlockattr_init()	-
	pthread_rwlockattr_destroy()	-
	pthread_rwlockattr_getpshared()	-
	pthread_rwlockattr_setpshared()	-
	<hr/>	
Functions Related to Semaphores	POSIX (libpthread)	Solaris (libthread)
	sem_init()	sema_init()
	sem_open()	-
	sem_close()	-
	sem_wait()	sema_wait()
	sem_trywait()	sema_trywait()
	sem_post()	sema_post()
	sem_getvalue()	-
	sem_unlink()	-
	sem_destroy()	sema_destroy()
	<hr/>	
Functions Related to fork() Clean Up	POSIX (libpthread)	Solaris (libthread)
	pthread_atfork()	-
	<hr/>	
Functions Related to Limits	POSIX (libpthread)	Solaris (libthread)
	pthread_once()	-
	<hr/>	
Functions Related to Debugging	POSIX (libpthread)	Solaris (libthread)
	-	thr_stksegment()
	<hr/>	

threads(3THR)

Synchronization	<p>POSIX (libpthread) Solaris (libthread) Multi-threaded behavior is asynchronous, and therefore, optimized for concurrent and parallel processing. As threads, always from within the same process and sometimes from multiple processes, share global data with each other, they are not guaranteed exclusive access to the shared data at any point in time. Securing mutually exclusive access to shared data requires synchronization among the threads. Both POSIX and Solaris implement four synchronization mechanisms: mutexes, condition variables, reader/writer locking (<i>optimized frequent-read occasional-write mutex</i>), and semaphores.</p> <p>Synchronizing multiple threads diminishes their concurrency. The coarser the grain of synchronization, that is, the larger the block of code that is locked, the lesser the concurrency.</p>						
MT fork()	<p>If a POSIX threads program calls <code>fork(2)</code>, it implicitly calls <code>fork1(2)</code>, which replicates only the calling thread. Should there be any outstanding mutexes throughout the process, the application should call <code>pthread_atfork(3THR)</code>, to wait for and acquire those mutexes, prior to calling <code>fork()</code>.</p>						
POSIX	<p>Scheduling allocation size per thread is greater than one. POSIX supports the following three scheduling policies:</p> <table><tr><td><code>SCHED_OTHER</code></td><td>Timesharing (TS) scheduling policy. It is based on the timesharing scheduling class.</td></tr><tr><td><code>SCHED_FIFO</code></td><td>First-In-First-Out (FIFO) scheduling policy. Threads scheduled to this policy, if not pre-empted by a higher priority, will proceed until completion. Threads whose contention scope is system (<code>PTHREAD_SCOPE_SYSTEM</code>) are in real-time (RT) scheduling class. The calling process must have a effective user ID of 0. <code>SCHED_FIFO</code> for threads whose contention scope's process (<code>PTHREAD_SCOPE_PROCESS</code>) is based on the TS scheduling class.</td></tr><tr><td><code>SCHED_RR</code></td><td>Round-Robin scheduling policy. Threads scheduled to this policy, if not pre-empted by a higher priority, will execute for a time period determined by the system. Threads whose contention scope is system (<code>PTHREAD_SCOPE_SYSTEM</code>) are in real-time (RT) scheduling class and the calling process must have a effective user ID of 0. <code>SCHED_RR</code> for threads whose contention scope is process (<code>PTHREAD_SCOPE_PROCESS</code>) is based on the TS scheduling class.</td></tr></table>	<code>SCHED_OTHER</code>	Timesharing (TS) scheduling policy. It is based on the timesharing scheduling class.	<code>SCHED_FIFO</code>	First-In-First-Out (FIFO) scheduling policy. Threads scheduled to this policy, if not pre-empted by a higher priority, will proceed until completion. Threads whose contention scope is system (<code>PTHREAD_SCOPE_SYSTEM</code>) are in real-time (RT) scheduling class. The calling process must have a effective user ID of 0. <code>SCHED_FIFO</code> for threads whose contention scope's process (<code>PTHREAD_SCOPE_PROCESS</code>) is based on the TS scheduling class.	<code>SCHED_RR</code>	Round-Robin scheduling policy. Threads scheduled to this policy, if not pre-empted by a higher priority, will execute for a time period determined by the system. Threads whose contention scope is system (<code>PTHREAD_SCOPE_SYSTEM</code>) are in real-time (RT) scheduling class and the calling process must have a effective user ID of 0. <code>SCHED_RR</code> for threads whose contention scope is process (<code>PTHREAD_SCOPE_PROCESS</code>) is based on the TS scheduling class.
<code>SCHED_OTHER</code>	Timesharing (TS) scheduling policy. It is based on the timesharing scheduling class.						
<code>SCHED_FIFO</code>	First-In-First-Out (FIFO) scheduling policy. Threads scheduled to this policy, if not pre-empted by a higher priority, will proceed until completion. Threads whose contention scope is system (<code>PTHREAD_SCOPE_SYSTEM</code>) are in real-time (RT) scheduling class. The calling process must have a effective user ID of 0. <code>SCHED_FIFO</code> for threads whose contention scope's process (<code>PTHREAD_SCOPE_PROCESS</code>) is based on the TS scheduling class.						
<code>SCHED_RR</code>	Round-Robin scheduling policy. Threads scheduled to this policy, if not pre-empted by a higher priority, will execute for a time period determined by the system. Threads whose contention scope is system (<code>PTHREAD_SCOPE_SYSTEM</code>) are in real-time (RT) scheduling class and the calling process must have a effective user ID of 0. <code>SCHED_RR</code> for threads whose contention scope is process (<code>PTHREAD_SCOPE_PROCESS</code>) is based on the TS scheduling class.						
Solaris	<p>Only scheduling policy supported is <code>SCHED_OTHER</code>, which is timesharing, based on the TS scheduling class.</p>						
ERRORS	<p>In a multi-threaded application, linked with <code>libpthread</code> or <code>libthread</code>, <code>EINTR</code> may be returned whenever another thread calls <code>fork(2)</code>, which calls <code>fork1(2)</code> instead. To ensure proper library linking order, use this option, rather than <code>-lthread</code>, to link with <code>libthread</code>.</p>						

-mt compiler option The `-mt` compiler option compiles and links for multithreaded code. It compiles source files with `-D_REENTRANT` and augments the set of support libraries to include `-lthread` in the required order.

To ensure proper library linking order, use this option rather than `-lthread` to link with `libthread`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe, Fork 1-Safe

POSIX `/usr/include/pthread.h /lib/libpthread.* /lib/librt.*`

Solaris `/usr/include/thread.h /usr/include/sched.h /lib/libthread.*`

SEE ALSO `crle(1)`, `fork(2)`, `pthread_atfork(3THR)`, `pthread_create(3THR)`, `attributes(5)`, `standards(5)`

Linker and Libraries Guide

thr_exit(3THR)

NAME	thr_exit – terminate the calling thread
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> void thr_exit(void *status);</pre>
DESCRIPTION	<p>The <code>thr_exit()</code> function terminates the calling thread, in a similar way that <code>exit(3C)</code> terminates the calling process. If the calling thread is not detached, then the thread's ID and the exit status specified by <i>status</i> are retained. The value <i>status</i> is then made available to any successful join with the terminating thread (see <code>thr_join(3THR)</code>); otherwise, <i>status</i> is disregarded allowing the thread's ID to be reclaimed immediately.</p> <p>Any cancellation cleanup handlers that have been pushed and not yet popped are popped in the reverse order that they were pushed and then executed. After all cancellation cleanup handlers have been executed, if the thread has any thread-specific data, appropriate destructor functions will be called in an unspecified order. Thread termination does not release any application visible process resources, including, but not limited to, mutexes and file descriptors, nor does it perform any process level cleanup actions, including, but not limited to, calling any <code>atexit()</code> routines that may exist.</p> <p>If any thread, including the <code>main()</code> thread, calls <code>thr_exit()</code>, only that thread will exit.</p> <p>If <code>main()</code> returns or exits (either implicitly or explicitly), or any thread explicitly calls <code>exit()</code>, the entire process will exit.</p> <p>The behavior of <code>thr_exit()</code> is undefined if called from a cancellation cleanup handler or destructor function that was invoked as a result of either an implicit or explicit call to <code>thr_exit()</code>.</p> <p>After a thread has terminated, the result of access to local (auto) variables of the thread is undefined. Thus, references to local variables of the exiting thread should not be used for the <code>thr_exit()</code> <i>status</i> parameter value.</p> <p>If any thread (except the <code>main()</code> thread) implicitly or explicitly returns, the result is the same as if the thread called <code>thr_exit()</code> and it will return the value of <i>status</i> as the exit code.</p> <p>The process will terminate with an exit status of 0 after the last non-daemon thread has terminated (including the <code>main()</code> thread). This behavior is the same as if the application had called <code>exit()</code> with a 0 argument at thread termination time.</p>
RETURN VALUES	The <code>thr_exit()</code> function cannot return to its caller.
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

thr_exit(3THR)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO exit(3C), thr_create(3THR), thr_join(3THR), thr_keycreate(3THR), attributes(5), standards(5)

NOTES Although only POSIX implements cancellation, cancellation can be used with Solaris threads, due to their interoperability.

The *status* argument should not reference any variables local to the calling thread.

thr_getconcurrency(3THR)

NAME	thr_getconcurrency, thr_setconcurrency – get or set thread concurrency level				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> int thr_setconcurrency(int new_level); int thr_getconcurrency(void);</pre>				
DESCRIPTION	<p>These functions are obsolete and maintained for compatibility only. The thr_setconcurrency() function updates the desired concurrency level that libthread maintains for the calling process. This value does not affect the behavior of the calling process.</p> <p>The thr_getconcurrency() function returns the current value for the desired concurrency level.</p>				
RETURN VALUES	<p>The thr_getconcurrency() function always returns the current value for the desired concurrency level.</p> <p>If successful, the thr_setconcurrency() function returns 0. Otherwise, a non-zero value is returned to indicate the error.</p>				
ERRORS	<p>The thr_setconcurrency() function will fail if:</p> <table><tr><td>EAGAIN</td><td>The specified concurrency level would cause a system resource to be exceeded.</td></tr><tr><td>EINVAL</td><td>The value for <i>new_level</i> is negative.</td></tr></table>	EAGAIN	The specified concurrency level would cause a system resource to be exceeded.	EINVAL	The value for <i>new_level</i> is negative.
EAGAIN	The specified concurrency level would cause a system resource to be exceeded.				
EINVAL	The value for <i>new_level</i> is negative.				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	thr_create(3THR), attributes(5), standards(5)				

NAME	thr_getprio, thr_setprio – access dynamic thread scheduling				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> int thr_setprio(thread_t <i>target_thread</i>, int <i>priority</i>); int thr_getprio(thread_t <i>target_thread</i>, int *<i>priority</i>);</pre>				
DESCRIPTION	<p>The <code>thr_setprio()</code> function dynamically changes the priority of the thread specified by <i>target_thread</i> within the current process to the priority specified by <i>priority</i>. By default, threads contend for synchronization objects based on fixed priorities that range from 0, the least significant, to 127. The <i>target_thread</i> will receive precedence by libthread over lower priority threads with respect to synchronization object contention.</p> <p>The <code>thr_getprio()</code> function stores the current priority for the thread specified by <i>target_thread</i> in the location pointed to by <i>priority</i>. Thread priorities regulate the order in which threads unblock from synchronization objects and are different from realtime priorities, which regulate and enforce access to CPU resources. Programs that need access to "real" priorities should use bound threads in the realtime class (see <code>prctl(2)</code>).</p>				
RETURN VALUES	If successful, the <code>thr_getprio()</code> and <code>thr_setprio()</code> return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>For each of the following conditions, these functions return an error number if the condition is detected.</p> <p>ESRCH The value specified by <i>target_thread</i> does not refer to an existing thread.</p> <p>The <code>thr_getprio()</code> and <code>thr_setprio()</code> functions may fail if:</p> <p>EINVAL The value of <i>priority</i> makes no sense for the scheduling class associated with the <i>target_thread</i>.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>prctl(2)</code> , <code>sched_setparam(3RT)</code> , <code>thr_create(3THR)</code> , <code>thr_suspend(3THR)</code> , <code>thr_yield(3THR)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

thr_join(3THR)

NAME	thr_join – wait for thread termination				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> int thr_join(thread_t <i>thread</i>, thread_t *<i>departed</i>, void **<i>status</i>);</pre>				
DESCRIPTION	<p>The <code>thr_join()</code> function suspends processing of the calling thread until the target <i>thread</i> completes. The <i>thread</i> argument must be a member of the current process and cannot be a detached thread. See <code>thr_create(3THR)</code>.</p> <p>If two or more threads wait for the same thread to complete, all will suspend processing until the thread has terminated, and then one thread will return successfully and the others will return with an error of <code>ESRCH</code>. The <code>thr_join()</code> function will not block processing of the calling thread if the target <i>thread</i> has already terminated.</p> <p>If a <code>thr_join()</code> call returns successfully with a non-null <i>status</i> argument, the value passed to <code>thr_exit(3THR)</code> by the terminating thread will be placed in the location referenced by <i>status</i>.</p> <p>If the target <i>thread</i> ID is 0, <code>thr_join()</code> finds and returns the status of a terminated undetached thread in the process. If no such thread exists, it suspends processing of the calling thread until a thread for which no other thread is waiting enters that state, at which time it returns successfully, or until all other threads in the process are either daemon threads or threads waiting in <code>thr_join()</code>, in which case it returns <code>EDEADLK</code>. See NOTES.</p> <p>If <i>departed</i> is not <code>NULL</code>, it points to a location that is set to the ID of the terminated thread if <code>thr_join()</code> returns successfully.</p>				
RETURN VALUES	If successful, <code>thr_join()</code> returns 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<table><tr><td><code>EDEADLK</code></td><td>A joining deadlock would occur, such as when a thread attempts to wait for itself, or the calling thread is waiting for any thread to exit and only daemon threads or waiting threads exist in the process.</td></tr><tr><td><code>ESRCH</code></td><td>No undetached thread could be found corresponding to the given thread ID.</td></tr></table>	<code>EDEADLK</code>	A joining deadlock would occur, such as when a thread attempts to wait for itself, or the calling thread is waiting for any thread to exit and only daemon threads or waiting threads exist in the process.	<code>ESRCH</code>	No undetached thread could be found corresponding to the given thread ID.
<code>EDEADLK</code>	A joining deadlock would occur, such as when a thread attempts to wait for itself, or the calling thread is waiting for any thread to exit and only daemon threads or waiting threads exist in the process.				
<code>ESRCH</code>	No undetached thread could be found corresponding to the given thread ID.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>wait(2)</code> , <code>thr_create(3THR)</code> , <code>thr_exit(3THR)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

NOTES | Using thr_join(3THR) in the following syntax,

```
while (thr_join(0, NULL, NULL) == 0);
```

will wait for the termination of all non-daemon threads, excluding threads that are themselves waiting in thr_join().

thr_keycreate(3THR)

NAME	thr_keycreate, thr_setspecific, thr_getspecific – thread-specific-data functions
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i>[<i>library...</i>] #include <thread.h> int thr_keycreate(thread_key_t *<i>keyp</i>, void (*<i>destructor</i>, void *<i>value</i>); int thr_setspecific(thread_key_t <i>key</i>, void *<i>value</i>); int thr_getspecific(thread_key_t <i>key</i>, void **<i>valuep</i>);</pre>
Create Key	<p>In general, thread key creation allocates a key that locates data specific to each thread in the process. The key is global to all threads in the process, which allows each thread to bind a value to the key once the key has been created. The key independently maintains specific values for each binding thread. The <code>thr_keycreate()</code> function allocates a global <i>key</i> namespace, pointed to by <i>keyp</i>, that is visible to all threads in the process. Each thread is initially bound to a private element of this <i>key</i>, which allows access to its thread-specific data.</p> <p>Upon key creation, a new key is assigned the value <code>NULL</code> for all active threads. Additionally, upon thread creation, all previously created keys in the new thread are assigned the value <code>NULL</code>.</p> <p>Optionally, a destructor function, <i>destructor</i>, may be associated with each <i>key</i>. Upon thread exit, if a <i>key</i> has a non-<code>NULL</code> <i>destructor</i> function and the thread has a non-<code>NULL</code> <i>value</i> associated with that <i>key</i>, the <i>destructor</i> function is called with the current associated <i>value</i>. If more than one <i>destructor</i> exists for a thread when it exits, the order of destructor calls is unspecified.</p>
Set Value	<p>Once a key has been created, each thread may bind a new <i>value</i> to the key using <code>thr_setspecific()</code>. The values are unique to the binding thread and are individually maintained. These values continue for the life of the calling thread.</p> <p>Proper synchronization of <i>key</i> storage and access must be ensured by the caller. The <i>value</i> argument to <code>thr_setspecific()</code> is generally a pointer to a block of dynamically allocated memory reserved by the calling thread for its own use. See <code>EXAMPLES</code>.</p> <p>At thread exit, the <i>destructor</i> function, which is associated at time of creation, is called and it uses the specific key value as its sole argument.</p>
Get Value	<code>thr_getspecific()</code> stores the current value bound to <i>key</i> for the calling thread into the location pointed to by <i>valuep</i> .
RETURN VALUES	If successful, <code>thr_keycreate()</code> , <code>thr_setspecific()</code> and <code>thr_getspecific()</code> return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	If the following conditions occur, <code>thr_keycreate()</code> returns the corresponding error number:

EAGAIN The system lacked the necessary resources to create another thread-specific data key.

ENOMEM Insufficient memory exists to create the key.

If the following conditions occur, `thr_keycreate()` and `thr_setspecific()` return the corresponding error number:

ENOMEM Insufficient memory exists to associate the value with the key.

The `thr_setspecific()` function returns the corresponding error number:

EINVAL The *key* value is invalid.

EXAMPLES

EXAMPLE 1 In this example, the thread-specific data in this function can be called from more than one thread without special initialization.

For each argument you pass to the executable of this example, a thread is created and privately bound to the string-value of that argument.

```
/* cc thisfile.c */

#include <thread.h>

#define _REENTRANT
void *thread_specific_data(), free();
#define MAX_ARGC 20
thread_t tid[MAX_ARGC];
int num_threads;

main( int argc, char *argv[] ) {
    int i;
    num_threads = argc - 1;
    for( i = 0; i < num_threads; i++)
        thr_create(NULL, 0, thread_specific_data, argv[i+1], 0, &tid[i]);
    for( i = 0; i < num_threads; i++)
        thr_join(tid[i], NULL, NULL);
} /* end main */

void *thread_specific_data(char private_data[])
{
    static mutex_t keylock; /* static ensures only one copy of keylock */
    static thread_key_t key;
    static int once_per_keyname = 0;
    void *tsd = NULL;

    if (!once_per_keyname) {
        mutex_lock(&keylock);
        if (!once_per_keyname) {
            thr_keycreate(&key, free);
            once_per_keyname++;
        }
        mutex_unlock(&keylock);
    }
    thr_getspecific(key, &tsd);
    if (tsd == NULL) {
```

thr_keycreate(3THR)

EXAMPLE 1 In this example, the thread-specific data in this function can be called from more than one thread without special initialization. *(Continued)*

```
    tsd = (void *)malloc(strlen(private_data) + 1);
    strcpy(tsd, private_data);
    thr_setspecific(key, tsd);
    thr_getspecific(key, &tsd);
    printf("tsd for %d = %s\n",thr_self(), tsd);
    thr_getspecific(key, &tsd);
    printf("tsd for %d remains %s\n",thr_self(), tsd);
    }
} /* end thread_specific_data */

void
free(void *v)  {
    /* application-specific clean-up function */
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO thr_exit(3THR), attributes(5), standards(5)

WARNINGS The thr_getspecific() and thr_getspecific() functions may be called either explicitly, or implicitly from a thread-specific data destructor function. Calling thr_setspecific() from a destructor may result in lost storage or infinite loops.

NAME thr_kill – send a signal to a thread

SYNOPSIS

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

#include <signal.h>
#include <thread.h>

int thr_kill(thread_t thread, int sig);
```

DESCRIPTION thr_kill() sends the *sig* signal to the thread designated by *thread*. *thread* must be a member of the same process as the calling thread. *sig* must be one of the signals listed in signal(3HEAD); with the exception of SIGLWP, SIGCANCEL, and SIGWAITING being reserved and off limits to thr_kill(). If *sig* is 0, a validity check is done for the existence of the target thread; no signal is sent.

RETURN VALUES Upon successful completion, thr_kill() returns 0. Otherwise, an error number is returned. In the event of failure, no signal is sent.

ERRORS

ESRCH	No thread was found that corresponded to the thread designated by <i>thread</i> ID.
EINVAL	The <i>sig</i> argument value is not zero and is an invalid or an unsupported signal number.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO kill(2), sigaction(2), raise(3C), thr_self(3THR), attributes(5), signal(3HEAD), standards(5)

thr_main(3THR)

- NAME** thr_main – identify the main thread
- SYNOPSIS**

```
cc -mt [ flag... ] file... [ library... ]  
  
#include <thread.h>  
  
int thr_main(void);
```
- DESCRIPTION** The thr_main() function returns one of the following:
- 1 if the calling thread is the main thread
 - 0 if the calling thread is not the main thread
 - 1 if libthread is not linked in or thread initialization has not completed
- FILES** /lib/libthread
- ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO thr_self(3THR), attributes(5)

NAME	thr_min_stack – return the minimum-allowable size for a thread’s stack				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> size_t thr_min_stack(void);</pre>				
DESCRIPTION	<p>When a thread is created with a user-supplied stack, the user must reserve enough space to run this thread. In a dynamically linked execution environment, it is very hard to know what the minimum stack requirements are for a thread. The function <code>thr_min_stack()</code> returns the amount of space needed to execute a null thread. This is a thread that was created to execute a null procedure. A thread that does something useful should have a stack size that is <code>thr_min_stack() + <some increment></code>.</p> <p>Most users should not be creating threads with user-supplied stacks. This functionality was provided to support applications that wanted complete control over their execution environment.</p> <p>Typically, users should let the threads library manage stack allocation. The threads library provides default stacks which should meet the requirements of any created thread.</p> <p><code>thr_min_stack()</code> will return the unsigned int <code>THR_MIN_STACK</code>, which is the minimum-allowable size for a thread’s stack.</p> <p>In this implementation the default size for a user-thread’s stack is one mega-byte. If the second argument to <code>thr_create(3THR)</code> is <code>NULL</code>, then the default stack size for the newly-created thread will be used. Otherwise, you may specify a stack-size that is at least <code>THR_MIN_STACK</code>, yet less than the size of your machine’s virtual memory.</p> <p>It is recommended that the default stack size be used.</p> <p>To determine the smallest-allowable size for a thread’s stack, execute the following:</p> <pre>/* cc thisfile.c -lthread */ #define _REENTRANT #include <thread.h> #include <stdio.h> main() { printf("thr_min_stack() returns %u\n",thr_min_stack()); }</pre>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>attributes(5)</code> , <code>standards(5)</code>				

thr_self(3THR)

NAME thr_self – get calling thread’s ID

SYNOPSIS

```
cc -mt [ flag... ] file...[ library... ]
#include <thread.h>

thread_t thr_self(void);
typedef(unsigned int thread_t);
```

DESCRIPTION thr_self() returns the thread ID of the calling thread.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO thr_create(3THR), attributes(5), standards(5)

NAME	thr_sigsetmask – change or examine calling thread’s signal mask						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> [<i>library...</i>] #include <thread.h> #include <signal.h> int thr_sigsetmask(int <i>how</i>, const sigset_t *<i>set</i>, sigset_t *<i>oset</i>);</pre>						
DESCRIPTION	<p>The <code>thr_sigsetmask()</code> function changes or examines a calling thread’s signal mask. Each thread has its own signal mask. A new thread inherits the calling thread’s signal mask and priority; however, pending signals are not inherited. Signals pending for a new thread will be empty.</p> <p>If the value of the argument <code>set</code> is not <code>NULL</code>, <code>set</code> points to a set of signals that can modify the currently blocked set. If the value of <code>set</code> is <code>NULL</code>, the value of <code>how</code> is insignificant and the thread’s signal mask is unmodified; thus, <code>thr_sigsetmask()</code> can be used to inquire about the currently blocked signals.</p> <p>The value of the argument <code>how</code> specifies the method in which the set is changed and takes one of the following values:</p> <table border="0"> <tr> <td style="padding-right: 20px;">SIG_BLOCK</td> <td>set corresponds to a set of signals to block. They are added to the current signal mask.</td> </tr> <tr> <td>SIG_UNBLOCK</td> <td>set corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.</td> </tr> <tr> <td>SIG_SETMASK</td> <td>set corresponds to the new signal mask. The current signal mask is replaced by <code>set</code>.</td> </tr> </table> <p>If the value of <code>oset</code> is not <code>NULL</code>, it points to the location where the previous signal mask is stored.</p>	SIG_BLOCK	set corresponds to a set of signals to block. They are added to the current signal mask.	SIG_UNBLOCK	set corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.	SIG_SETMASK	set corresponds to the new signal mask. The current signal mask is replaced by <code>set</code> .
SIG_BLOCK	set corresponds to a set of signals to block. They are added to the current signal mask.						
SIG_UNBLOCK	set corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.						
SIG_SETMASK	set corresponds to the new signal mask. The current signal mask is replaced by <code>set</code> .						
RETURN VALUES	Upon successful completion, the <code>thr_sigsetmask()</code> function returns 0. Otherwise, it returns a non-zero value.						
ERRORS	<p>The <code>thr_sigsetmask()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The value of <code>how</code> is not defined and <code>oset</code> is <code>NULL</code>.</td> </tr> </table>	EINVAL	The value of <code>how</code> is not defined and <code>oset</code> is <code>NULL</code> .				
EINVAL	The value of <code>how</code> is not defined and <code>oset</code> is <code>NULL</code> .						
EXAMPLES	<p>EXAMPLE 1 The following example shows how to create a default thread that can serve as a signal catcher/handler with its own signal mask. <code>new</code> will have a different value from the creator’s signal mask.</p> <p>As POSIX threads and Solaris threads are fully compatible even within the same process, this example uses <code>pthread_create(3THR)</code> if you execute <code>a.out 0</code>, or <code>thr_create(3THR)</code> if you execute <code>a.out 1</code>.</p> <p>In this example:</p> <ul style="list-style-type: none"> ■ <code>sigemptyset(3C)</code> initializes a null signal set, <code>new</code>. <code>sigaddset(3C)</code> packs the signal, <code>SIGINT</code>, into that new set. 						

thr_sigsetmask(3THR)

EXAMPLE 1 The following example shows how to create a default thread that can serve as a signal catcher/handler with its own signal mask. `new` will have a different value from the creator's signal mask. (Continued)

- Either `pthread_sigmask()` or `thr_sigsetmask()` is used to mask the signal, `SIGINT` (CTRL-C), from the calling thread, which is `main()`. The signal is masked to guarantee that only the new thread will receive this signal.
- `pthread_create()` or `thr_create()` creates the signal-handling thread.
- Using `pthread_join(3THR)` or `thr_join(3THR)`, `main()` then waits for the termination of that signal-handling thread, whose ID number is `user_threadID`; after which, `main()` will `sleep(3C)` for 2 seconds, and then the program terminates.
- The signal-handling thread, handler:
 - Assigns the handler `interrupt()` to handle the signal `SIGINT`, by the call to `sigaction(2)`.
 - Resets its own signal set to *not block* the signal, `SIGINT`.
 - Sleeps for 8 seconds to allow time for the user to deliver the signal, `SIGINT`, by pressing the CTRL-C.

```
/* cc thisfile.c -lthread -lpthread */
#define _REENTRANT /* basic first 3-lines for threads */
#include <pthread.h>
#include <thread.h>

thread_t user_threadID;
sigset_t new;
void *handler( ), interrupt( );

main( int argc, char *argv[ ] ){
    test_argv(argv[1]);

    sigemptyset(&new);
    sigaddset(&new, SIGINT);
    switch(*argv[1]) {

        case '0': /* POSIX */
            pthread_sigmask(SIG_BLOCK, &new, NULL);
            pthread_create(&user_threadID, NULL, handler, argv[1]);
            pthread_join(user_threadID, NULL);
            break;

        case '1': /* Solaris */
            thr_sigsetmask(SIG_BLOCK, &new, NULL);
            thr_create(NULL, 0, handler, argv[1], 0, &user_threadID);
            thr_join(user_threadID, NULL, NULL);
            break;
    } /* switch */

    printf("thread handler, # %d, has exited\n",user_threadID);
    sleep(2);
    printf("main thread, # %d is done\n", thr_self( ));
}
```

EXAMPLE 1 The following example shows how to create a default thread that can serve as a signal catcher/handler with its own signal mask. new will have a different value from the creator's signal mask. (Continued)

```

} /* end main */

struct sigaction act;

void *
handler(char argv1[ ])
{
    act.sa_handler = interrupt;
    sigaction(SIGINT, &act, NULL);
    switch(*argv1){
        case '0': /* POSIX */
            pthread_sigmask(SIG_UNBLOCK, &new, NULL);
            break;
        case '1': /* Solaris */
            thr_sigsetmask(SIG_UNBLOCK, &new, NULL);
            break;
    }
    printf("\n Press CTRL-C to deliver SIGINT signal to the process\n");
    sleep(8); /* give user time to hit CTRL-C */
}

void
interrupt(int sig)
{
    printf("thread %d caught signal %d\n", thr_self( ), sig);
}

void test_argv(char argv1[ ]) {
    if(argv1 == NULL) {
        printf("use 0 as arg1 to use thr_create( );\n \
or use 1 as arg1 to use pthread_create( )\n");
        exit(NULL);
    }
}

```

EXAMPLE 2

In the last example, the handler thread served as a signal-handler while also taking care of activity of its own (in this case, sleeping, although it could have been some other activity). A thread could be completely dedicated to signal-handling simply by waiting for the delivery of a selected signal by blocking with `sigwait(2)`. The two subroutines in the previous example, `handler()` and `interrupt()`, could have been replaced with the following routine:

```

void *
handler( )
{ int signal;
  printf("thread %d waiting for you to press the CTRL-C keys\n",
        thr_self( ));
  sigwait(&new, &signal);
  printf("thread %d has received the signal %d \n", thr_self( ), signal);
}

```

thr_sigsetmask(3THR)

EXAMPLE 2 (Continued)

```
}
/*pthread_create( ) and thr_create( ) would use NULL instead of
   argv[1] for the arg passed to handler( ) */
```

In this routine, one thread is dedicated to catching and handling the signal specified by the set new, which allows `main()` and all of its other sub-threads, created *after* `pthread_sigmask()` or `thr_sigsetmask()` masked that signal, to continue uninterrupted. Any use of `sigwait(2)` should be such that all threads block the signals passed to `sigwait(2)` at all times. Only the thread that calls `sigwait()` will get the signals. The call to `sigwait(2)` takes two arguments.

For this type of background dedicated signal-handling routine, you may wish to use a Solaris daemon thread by passing the argument `THR_DAEMON` to `thr_create()`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe and Async-Signal-Safe

SEE ALSO `sigaction(2)`, `sigprocmask(2)`, `sigwait(2)`, `cond_wait(3THR)`, `pthread_create(3THR)`, `pthread_join(3THR)`, `pthread_self(3THR)`, `sigsetops(3C)`, `sleep(3C)`, `attributes(5)`, `standards(5)`

NOTES It is not possible to block signals that cannot be ignored (see `sigaction(2)`). If using the threads library, it is not possible to block or unblock the signals `SIGWAITING`, `SIGLWP`, or `SIGCANCEL`. This restriction is quietly enforced by the threads library.

Using `sigwait(2)` in a dedicated thread allows asynchronously generated signals to be managed synchronously; however, `sigwait(2)` should never be used to manage synchronously generated signals.

Synchronously generated signals are exceptions that are generated by a thread and are directed at the thread causing the exception. Since `sigwait()` blocks waiting for signals, the blocking thread cannot receive a synchronously generated signal.

If `sigprocmask(2)` is used in a multi-threaded program, it will be the same as if `thr_sigsetmask()` or `pthread_sigmask()` has been called. POSIX leaves the semantics of the call to `sigprocmask(2)` unspecified in a multi-threaded process, so programs that care about POSIX portability should not depend on this semantic.

If a signal is delivered while a thread is waiting on a condition variable, the `cond_wait(3THR)` function will be interrupted and the handler will be executed. The state of the lock protecting the condition variable is undefined while the thread is executing the signal handler.

`thr_sigsetmask(3THR)`

Signals that are generated synchronously should not be masked. If such a signal is blocked and delivered, the receiving process is killed.

thr_stksegment(3THR)

NAME	thr_stksegment – get thread stack address and size				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i>[<i>library...</i>] #include <thread.h> #include <signal.h> int thr_stksegment(stack_t *ss);</pre>				
DESCRIPTION	<p>The <code>thr_stksegment()</code> function returns, in its <code>stack_t</code> argument, the address and size of the calling thread's stack.</p> <p>The <code>stack_t</code> structure includes the following members:</p> <pre>void *ss_sp size_t ss_size int ss_flags</pre> <p>On successful return from <code>thr_stksegment()</code>, <code>ss_sp</code> contains the high address of the caller's stack and <code>ss_size</code> contains the size of the stack in bytes. The <code>ss_flags</code> member is always 0. Note that the meaning of <code>ss_sp</code> is reversed from other uses of <code>stack_t</code> such as <code>sigaltstack(2)</code> where <code>ss_sp</code> is the low address.</p> <p>The stack information provided by <code>thr_stksegment()</code> is typically used by debuggers, garbage collectors, and similar applications. Most applications should not require such information.</p>				
RETURN VALUES	The <code>thr_stksegment()</code> function returns 0 if the thread stack address and size were successfully retrieved. Otherwise, it returns a non-zero error value.				
ERRORS	The <code>thr_stksegment()</code> function will fail if: EAGAIN The stack information for the thread is not available because the thread's initialization is not yet complete, or the thread is an internal thread.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>sigaltstack(2)</code> , <code>thr_create(3THR)</code> , <code>attributes(5)</code>				

NAME thr_suspend, thr_continue – suspend or continue thread execution

SYNOPSIS

```
cc -mt [ flag... ] file...[ library... ]
#include <thread.h>

int thr_suspend(thread_t target_thread) ;
int thr_continue(thread_t target_thread) ;
```

DESCRIPTION

The `thr_suspend()` function immediately suspends the execution of the thread specified by `target_thread`. On successful return from `thr_suspend()`, the suspended thread is no longer executing. Once a thread is suspended, subsequent calls to `thr_suspend()` have no effect.

The `thr_continue()` function resumes the execution of a suspended thread. Once a suspended thread is continued, subsequent calls to `thr_continue()` have no effect.

A suspended thread will not be awakened by a signal. The signal stays pending until the execution of the thread is resumed by `thr_continue()`.

RETURN VALUES If successful, the `thr_suspend()` and `thr_continue()` functions return 0. Otherwise, a non-zero value is returned to indicate the error.

ERRORS The `thr_suspend()` or `thr_continue()` functions will fail if:

ESRCH	The <code>target_thread</code> cannot be found in the current process.
ECANCELED	The <code>target_thread</code> was not suspended because a subsequent <code>thr_continue()</code> occurred before the suspend completed.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `thr_create(3THR)`, `thr_join(3THR)`, `attributes(5)`, `standards(5)`

thr_yield(3THR)

NAME thr_yield – yield to another thread

SYNOPSIS cc -mt [*flag...*] *file...*[*library...*]

```
#include <thread.h>
void thr_yield(void);
```

DESCRIPTION The `thr_yield()` function causes the current thread to yield its execution in favor of another thread with the same or greater priority.

RETURN VALUES The `thr_yield()` function returns nothing and does not set `errno`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `thr_setprio(3THR)`, `attributes(5)`, `standards(5)`

NAME	timer_create – create a timer								
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <signal.h> #include <time.h> int timer_create(clockid_t clock_id, struct sigevent *evp, timer_t *timerid);</pre>								
DESCRIPTION	<p>The <code>timer_create()</code> function creates a timer using the specified clock, <code>clock_id</code>, as the timing base. The <code>timer_create()</code> function returns, in the location referenced by <code>timerid</code>, a timer ID of type <code>timer_t</code> used to identify the timer in timer requests. This timer ID will be unique within the calling process until the timer is deleted. The particular clock, <code>clock_id</code>, is defined in <code><time.h></code>. The timer whose ID is returned will be in a disarmed state upon return from <code>timer_create()</code>.</p> <p>The <code>evp</code> argument, if non-null, points to a <code>sigevent</code> structure. This structure, allocated by the application, defines the asynchronous notification that will occur when the timer expires. If the <code>evp</code> argument is <code>NULL</code>, the effect is as if the <code>evp</code> argument pointed to a <code>sigevent</code> structure with the <code>sigev_notify</code> member having the value <code>SIGEV_SIGNAL</code>, the <code>sigev_signo</code> having a default signal number, and the <code>sigev_value</code> member having the value of the timer ID, <code>timerid</code>.</p> <p>The system defines a set of clocks that can be used as timing bases for per-process timers. The following values for <code>clock_id</code> are supported:</p> <table border="0" style="margin-left: 20px;"> <tr> <td><code>CLOCK_REALTIME</code></td> <td>wall clock</td> </tr> <tr> <td><code>CLOCK_VIRTUAL</code></td> <td>user CPU usage clock</td> </tr> <tr> <td><code>CLOCK_PROF</code></td> <td>user and system CPU usage clock</td> </tr> <tr> <td><code>CLOCK_HIGHRES</code></td> <td>non-adjustable, high-resolution clock</td> </tr> </table> <p>For timers created with a <code>clock_id</code> of <code>CLOCK_HIGHRES</code>, the system will attempt to use an optimal hardware source. This may include, but is not limited to, per-CPU timer sources. The actual hardware source used is transparent to the user and may change over the lifetime of the timer. For example, if the caller that created the timer were to change its processor binding or its processor set, the system may elect to drive the timer with a hardware source that better reflects the new binding. Timers based on a <code>clock_id</code> of <code>CLOCK_HIGHRES</code> are ideally suited for interval timers that have minimal jitter tolerance.</p> <p>Timers are not inherited by a child process across a <code>fork(2)</code> and are disarmed and deleted by a call to one of the <code>exec</code> functions (see <code>exec(2)</code>).</p>	<code>CLOCK_REALTIME</code>	wall clock	<code>CLOCK_VIRTUAL</code>	user CPU usage clock	<code>CLOCK_PROF</code>	user and system CPU usage clock	<code>CLOCK_HIGHRES</code>	non-adjustable, high-resolution clock
<code>CLOCK_REALTIME</code>	wall clock								
<code>CLOCK_VIRTUAL</code>	user CPU usage clock								
<code>CLOCK_PROF</code>	user and system CPU usage clock								
<code>CLOCK_HIGHRES</code>	non-adjustable, high-resolution clock								
RETURN VALUES	<p>Upon successful completion, <code>timer_create()</code> returns 0 and updates the location referenced by <code>timerid</code> to a <code>timer_t</code>, which can be passed to the per-process timer calls. If an error occurs, the function returns <code>-1</code> and sets <code>errno</code> to indicate the error. The value of <code>timerid</code> is undefined if an error occurs.</p>								
ERRORS	<p>The <code>timer_create()</code> function will fail if:</p>								

timer_create(3RT)

EAGAIN	The system lacks sufficient signal queuing resources to honor the request, or the calling process has already created all of the timers it is allowed by the system.
EINVAL	The specified clock ID, <i>clock_id</i> , is not defined.
ENOSYS	The <code>timer_create()</code> function is not supported by the system.
EPERM	The specified clock ID, <i>clock_id</i> , is <code>CLOCK_HIGHRES</code> and the effective user of the caller is not superuser.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe with exceptions

SEE ALSO `exec(2)`, `fork(2)`, `time(2)`, `clock_gettime(3RT)`, `signal(3C)`, `timer_delete(3RT)`, `timer_settime(3RT)`, `attributes(5)`

NAME timer_delete – delete a timer

SYNOPSIS

```
cc [ flag... ] file... -lrt [ library... ]
#include <time.h>

int timer_delete(timer_t timerid);
```

DESCRIPTION The `timer_delete()` function deletes the specified timer, *timerid*, previously created by the `timer_create(3RT)` function. If the timer is armed when `timer_delete()` is called, the behavior will be as if the timer is automatically disarmed before removal. The disposition of pending signals for the deleted timer is unspecified.

RETURN VALUES If successful, the function returns 0. Otherwise, the function returns -1 and sets `errno` to indicate the error.

ERRORS The `timer_delete()` function will fail if:

`EINVAL` The timer ID specified by *timerid* is not a valid timer ID.

`ENOSYS` The `timer_delete()` function is not supported by the system.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe with exceptions

SEE ALSO `timer_create(3RT)`, `attributes(5)`

timer_settime(3RT)

NAME	timer_settime, timer_gettime, timer_getoverrun – per-process timers
SYNOPSIS	<pre>cc [flag...] file... -lrt [library...] #include <time.h> int timer_settime(timer_t timerid, int flags, const struct itimerspec *value, struct itimerspec *ovalue); int timer_gettime(timer_t timerid, struct itimerspec *value); int timer_getoverrun(timer_t timerid);</pre>
DESCRIPTION	<p>The <code>timer_settime()</code> function sets the time until the next expiration of the timer specified by <code>timerid</code> from the <code>it_value</code> member of the <code>value</code> argument and arm the timer if the <code>it_value</code> member of <code>value</code> is non-zero. If the specified timer was already armed when <code>timer_settime()</code> is called, this call resets the time until next expiration to the <code>value</code> specified. If the <code>it_value</code> member of <code>value</code> is 0, the timer is disarmed. The effect of disarming or resetting a timer on pending expiration notifications is unspecified.</p> <p>If the flag <code>TIMER_ABSTIME</code> is not set in the argument <code>flags</code>, <code>timer_settime()</code> behaves as if the time until next expiration is set to be equal to the interval specified by the <code>it_value</code> member of <code>value</code>. That is, the timer expires in <code>it_value</code> nanoseconds from when the call is made. If the flag <code>TIMER_ABSTIME</code> is set in the argument <code>flags</code>, <code>timer_settime()</code> behaves as if the time until next expiration is set to be equal to the difference between the absolute time specified by the <code>it_value</code> member of <code>value</code> and the current value of the clock associated with <code>timerid</code>. That is, the timer expires when the clock reaches the value specified by the <code>it_value</code> member of <code>value</code>. If the specified time has already passed, the function succeeds and the expiration notification is made.</p> <p>The reload value of the timer is set to the value specified by the <code>it_interval</code> member of <code>value</code>. When a timer is armed with a non-zero <code>it_interval</code>, a periodic (or repetitive) timer is specified.</p> <p>Time values that are between two consecutive non-negative integer multiples of the resolution of the specified timer will be rounded up to the larger multiple of the resolution. Quantization error will not cause the timer to expire earlier than the rounded time value.</p> <p>If the argument <code>ovalue</code> is not <code>NULL</code>, the function <code>timer_settime()</code> stores, in the location referenced by <code>ovalue</code>, a value representing the previous amount of time before the timer would have expired or 0 if the timer was disarmed, together with the previous timer reload value. The members of <code>ovalue</code> are subject to the resolution of the timer, and they are the same values that would be returned by a <code>timer_gettime()</code> call at that point in time.</p> <p>The <code>timer_gettime()</code> function stores the amount of time until the specified timer, <code>timerid</code>, expires and the reload value of the timer into the space pointed to by the <code>value</code> argument. The <code>it_value</code> member of this structure contains the amount of time before</p>

timer_settime(3RT)

the timer expires, or 0 if the timer is disarmed. This value is returned as the interval until timer expiration, even if the timer was armed with absolute time. The `it_interval` member of *value* contains the reload value last set by `timer_settime()`.

Only a single signal will be queued to the process for a given timer at any point in time. When a timer for which a signal is still pending expires, no signal will be queued, and a timer overrun occurs. When a timer expiration signal is delivered to or accepted by a process, the `timer_getoverrun()` function returns the timer expiration overrun count for the specified timer. The overrun count returned contains the number of extra timer expirations that occurred between the time the signal was generated (queued) and when it was delivered or accepted, up to but not including an implementation-dependent maximum of `DELAYTIMER_MAX`. If the number of such extra expirations is greater than or equal to `DELAYTIMER_MAX`, then the overrun count will be set to `DELAYTIMER_MAX`. The value returned by `timer_getoverrun()` applies to the most recent expiration signal delivery or acceptance for the timer. If no expiration signal has been delivered for the timer, the meaning of the overrun count returned is undefined.

RETURN VALUES

If the `timer_settime()` or `timer_gettime()` functions succeed, 0 is returned. If an error occurs for either of these functions, -1 is returned, and `errno` is set to indicate the error. If the `timer_getoverrun()` function succeeds, it returns the timer expiration overrun count as explained above.

ERRORS

The `timer_settime()`, `timer_gettime()` and `timer_getoverrun()` functions will fail if:

- EINVAL The *timerid* argument does not correspond to a timer returned by `timer_create(3RT)` but not yet deleted by `timer_delete(3RT)`.
- ENOSYS The `timer_settime()`, `timer_gettime()`, and `timer_getoverrun()` functions are not supported by the system. The `timer_settime()` function will fail if:
- EINVAL A *value* structure specified a nanosecond value less than zero or greater than or equal to 1000 million.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`clock_settime(3RT)`, `timer_create(3RT)`, `timer_delete(3RT)`, `attributes(5)`, `time(3HEAD)`

timer_settime(3RT)

Index

Numbers and Symbols

get a synchronization object handle from a synchronization object's address —
td_ta_map_addr2sync, 254

A

access dynamic thread scheduling

— thr_getprio, 289

— thr_setprio, 289

access dynamic thread scheduling parameters

— pthread_getschedparam, 150

— pthread_setschedparam, 150

aiocancel — cancel an asynchronous operation, 16

aio_cancel — cancel asynchronous I/O request, 17

aio_fsync — asynchronous file synchronization, 21

aio_read — asynchronous read and write operations, 25

aioread — read or write asynchronous I/O operations, 23

aio_return — retrieve return status of asynchronous I/O operation, 28

aio_suspend — wait for asynchronous I/O request, 29

aiowait — wait for completion of asynchronous I/O operation, 31

aio_write — asynchronous write to a file, 32

aiowrite — read or write asynchronous I/O operations, 23

allocate and deallocate process handles for libthread_db

— td_ta_delete, 256

— td_ta_get_ph, 256

— td_ta_new, 256

asynchronous file synchronization

— aio_sync, 21

asynchronous I/O

— aio_cancel, 17

— aiocancel, 16

— aiowait, 31

retrieve return status — aio_return, 28

asynchronous read and write operations

— aio_read, aio_write, 25

asynchronous write to a file — aio_write, 32

B

bind or unbind the current thread with the door server pool

— door_bind, 50

— door_unbind, 50

C

cancellation — overview of concepts related to POSIX thread cancellation, 35

Cancellation, 35

cancellation — overview of concepts related to POSIX thread cancellation (*continued*)

- Cancellation Points, 36
- Cancellation State, 37
- Cancellation Type, 37
- Cancel-Safe, 38
- Cleanup Handlers, 36
- Planning Steps, 35
- POSIX Threads Only, 38

change or examine calling thread's signal mask — pthread_sigmask, 195, 299

change the priority ceiling of a mutex

- pthread_mutex_getprioceiling, 173
- pthread_mutex_setprioceiling, 173

clock_getres — high-resolution clock operations, 41

clock_gettime — high-resolution clock operations, 41

clock_settime — high-resolution clock operations, 41

collect target process statistics for libthread_db

- td_ta_enable_stats, 247
- td_ta_get_stats, 247
- td_ta_reset_stats, 247

compare thread IDs — pthread_equal, 146

concepts related to POSIX pthreads and Solaris threads and the libpthread and libthread libraries — threads, 279

concepts related to condition variables — condition, 48

concepts relating to mutual exclusion locks — mutex, 90

cond_broadcast — condition variables, 43

cond_destroy — condition variables, 43

cond_init — condition variables, 43

- Condition Signaling, 44
- Condition Wait, 44
- Destroy, 45
- Initialize, 43

condition variables — cond_init, 43, 48

condition — concepts related to condition variables

- Condition Signaling, 49
- Condition Wait, 48
- Destroy, 49
- Initialize, 48

cond_reltimedwait — condition variables, 43

cond_signal — condition variables, 43

cond_timedwait — condition variables, 43

cond_wait — condition variables, 43

convert a thread id or thread address to a thread handle

- td_ta_map_addr2thr, 255
- td_ta_map_id2thr, 255

create a door descriptor — door_create, 56

create a thread — thr_create, 273

create a thread — pthread_create, 141

create cancellation point in the calling thread. — pthread_testcancel, 200

create thread-specific data key — pthread_key_create, 155

D

delete thread-specific data key — pthread_key_delete, 157

detach a thread — pthread_detach, 145

door_bind — bind or unbind the current thread with the door server pool, 50

door_call — invoke the function associated with a door descriptor, 53

door_create — create a door descriptor, 56

door_cred — return credential information associated with the client, 58

door_info — return information associated with a door descriptor, 59

door_return — return from a door invocation, 61

door_revoke — revoke access to a door descriptor, 62

door_server_create — specify an alternative door server thread creation function, 63

door_unbind — bind or unbind the current thread with the door server pool, 50

E

enable or disable cancellation — pthread_setcancelstate, 191

enabling or disabling cancellation — pthread_setcancelstate, 191

F

`fdatasync` — synchronize a file's data, 65

G

get a thread's thread-specific data for
`libthread_db` library of interfaces —
`td_thr_tsd`, 271

get and set prioceiling attribute of mutex
attribute object
— `pthread_mutexattr_getprioceiling`, 159
— `pthread_mutexattr_setprioceiling`, 159

get and set process-shared attribute
— `pthread_mutexattr_getpshared`, 164
— `pthread_mutexattr_setpshared`, 164

get and set process-shared attribute of
read-write lock attributes object
— `pthread_rwlockattr_getpshared`, 181
— `pthread_rwlockattr_setpshared`, 181

get and set protocol attribute of mutex attribute
object
— `pthread_mutexattr_getprotocol`, 161
— `pthread_mutexattr_setprotocol`, 161

get calling thread's ID — `pthread_self`, 190,
298

get execution time limits —
`sched_rr_get_interval`, 209

get message queue attributes — `mq_getattr`, 78

get or set a mutex type
— `pthread_mutexattr_gettype`, 168
— `pthread_mutexattr_settype`, 168

get or set contentionscope attribute
— `pthread_attr_getscope`, 122
— `pthread_attr_setscope`, 122

get or set detachstate attribute
— `pthread_attr_getdetachstate`, 115
— `pthread_attr_setdetachstate`, 115

get or set inheritsched attribute
— `pthread_attr_getinheritsched`, 118
— `pthread_attr_setinheritsched`, 118

get or set level of concurrency
— `pthread_getconcurrency`, 148
— `pthread_setconcurrency`, 148

get or set schedparam attribute
— `pthread_attr_getschedparam`, 120
— `pthread_attr_setschedparam`, 120

get or set schedpolicy attribute
— `pthread_attr_getschedpolicy`, 121
— `pthread_attr_setschedpolicy`, 121

get or set stackaddr attribute
— `pthread_attr_getstackaddr`, 123
— `pthread_attr_setstackaddr`, 123

get or set stacksize attribute
— `pthread_attr_getstacksize`, 124
— `pthread_attr_setstacksize`, 124

get or set the process-shared condition variable
attributes
— `pthread_condattr_getpshared`, 130
— `pthread_condattr_setpshared`, 130

get or set the thread guardsize attribute
— `pthread_attr_getguardsize`, 116
— `pthread_attr_setguardsize`, 116

get scheduling parameter limits
— `sched_get_priority_max`, 207
— `sched_get_priority_min`, 207

get scheduling parameters —
`sched_getparam`, 206

get scheduling policy —
`sched_getscheduler`, 208

get thread information in `libthread_db` library
of interfaces — `td_thr_get_info`, 264

get thread stack address and size —
`thr_stksegment`, 304

gets the total number of threads in a process for
`libthread_db` — `td_ta_get_nthreads`, 253

I

initialization function for `libthread_db` library of
interfaces — `td_init`, 241

initialize and destroy mutex attributes object
— `pthread_mutexattr_destroy`, 170
— `pthread_mutexattr_init`, 170

initialize and destroy read-write lock attributes
object
— `pthread_rwlockattr_destroy`, 182
— `pthread_rwlockattr_init`, 182

initialize and destroy threads attribute object
— `pthread_attr_destroy`, 125
— `pthread_attr_init`, 125

initialize dynamic package —
`pthread_once`, 180

- initialize or destroy a mutex
 - `pthread_mutex_destroy`, 175
 - `pthread_mutex_init`, 175
- initialize or destroy a read-write lock object
 - `pthread_rwlock_destroy`, 183
 - `pthread_rwlock_init`, 183
- initialize or destroy condition variable attributes object
 - `pthread_condattr_destroy`, 132
 - `pthread_condattr_init`, 132
- initialize or destroy condition variables
 - `pthread_cond_destroy`, 134
 - `pthread_cond_init`, 134
- interfaces in `libthread_db` that target process memory access
 - `ps_phread`, 110
 - `ps_pdwrite`, 110
 - `ps_ptread`, 110
 - `ps_ptwrite`, 110
- invoke the function associated with a door descriptor — `door_call`, 53
- I/O, asynchronous
 - cancel request — `aio_cancel`, 17
 - file synchronization — `aio_sync`, 21
 - retrieve return status — `aio_return`, 28
- I/O, requests
 - list — `lio_listio`, 73
- iterate over the set of locks owned by a thread
 - `td_thr_lockowner`, 267
- iterator functions on process handles from `libthread_db` — `td_ta_sync_iter`, 259

L

- library of interfaces for monitoring and manipulating threads-related aspects of multithreaded programs — `libthread_db`, 68
- `libthread_db` — library of interfaces for monitoring and manipulating threads-related aspects of multithreaded programs, 68
- `lio_listio` — list directed I/O, 73
- list directed I/O — `lio_listio`, 73
- lock or attempt to lock a read-write lock object for reading
 - `pthread_rwlock_rdlock`, 185

- lock or attempt to lock a read-write lock object for reading (*continued*)
 - `pthread_rwlock_tryrdlock`, 185
- lock or attempt to lock a read-write lock object for writing
 - `pthread_rwlock_trywrlock`, 188
 - `pthread_rwlock_wrlock`, 188
- lock or unlock a mutex
 - `pthread_mutex_lock`, 177
 - `pthread_mutex_trylock`, 177
 - `pthread_mutex_unlock`, 177
- looks up the symbol in the symbol table of the load object in the target process — `ps_pglobal_lookup`, 109

M

- make a mutex consistent after owner death — `pthread_mutex_consistent_np`, 171
- manage thread signals for `libthread_db`
 - `td_thr_setsigpending`, 269
 - `td_thr_sigsetmask`, 269
- manage thread-specific data
 - `pthread_getspecific`, 152
 - `pthread_setspecific`, 152
- memory object, shared
 - open — `shm_open`, 233
 - remove — `shm_unlink`, 236
- message queue
 - close — `mq_close`, 77
 - notify process (or thread) — `mq_notify`, 79
 - open — `mq_open`, 81
 - receive a message from — `mq_receive`, 84
 - remove — `mq_unlink`, 89
 - send message to — `mq_send`, 86
 - set attributes — `mq_setattr`, 88
- `mq_close` — close a message queue, 77
- `mq_getattr` — get message queue attributes, 78
- `mq_notify` — notify process (or thread) that a message is available on a queue, 79
- `mq_open` — open a message queue, 81
- `mq_receive` — receive a message from a message queue, 84
- `mq_send` — send a message to a message queue, 86

mq_setattr — set/get message queue attributes, 88

mq_unlink — remove a message queue, 89

mutex — concepts relating to mutual exclusion locks, 90

- Caveats, 91
- Initialization, 90

mutex_destroy — mutual exclusion locks, 92

mutex_init — mutual exclusion locks, 92

- Destroy, 95
- Dynamically Allocated Mutexes, 101
- Initialize, 92
- Interprocess Locking, 98
- Lock and Unlock, 94
- Multiple Instruction Single Data, 96
- Single Gate, 96
- Solaris Interprocess Robust Locking, 99

mutex_lock — mutual exclusion locks, 92

mutex_trylock — mutual exclusion locks, 92

mutex_unlock — mutual exclusion locks, 92

mutual exclusion locks

- mutex_destroy, 92
- mutex_init, 92
- mutex_lock, 92
- mutex_trylock, 92
- mutex_unlock, 92

N

nanosleep — high resolution sleep, 103

O

operations on a synchronization object in libthread_db — td_sync_get_info, 243

overview of concepts related to POSIX thread cancellation — cancellation, 35

P

placeholder for future logging functionality — td_log, 242

pop a thread cancellation cleanup handler — pthread_cleanup_pop, 128

preemption control

- schedctl_exit, 204
- schedctl_init, 204
- schedctl_lookup, 204
- schedctl_start, 204
- schedctl_stop, 204

process and LWP control in libthread_db

- ps_kill, 111
- ps_lcontinue, 111
- ps_lrolltoaddr, 111
- ps_lstop, 111
- ps_pcontinue, 111
- ps_pstop, 111

process service interfaces — proc_service, 105

proc_service — process service interfaces, 105

- IA, 104
- SPARC, 104

ps_kill — process and LWP control in libthread_db, 111

ps_lcontinue — process and LWP control in libthread_db, 111

ps_lgetfpregs — routines that access the target process register in libthread_db, 107

ps_lgetregs — routines that access the target process register in libthread_db, 107

ps_lgetxregs — routines that access the target process register in libthread_db, 107

ps_lgetxregsize — routines that access the target process register in libthread_db, 107

ps_lrolltoaddr — process and LWP control in libthread_db, 111

ps_lsetfpregs — routines that access the target process register in libthread_db, 107

ps_lsetregs — routines that access the target process register in libthread_db, 107

ps_lsetxregs — routines that access the target process register in libthread_db, 107

ps_lstop — process and LWP control in libthread_db, 111

ps_pcontinue — process and LWP control in libthread_db, 111

ps_pdread — interfaces in libthread_db that target process memory access, 110

ps_pdwrites — interfaces in libthread_db that target process memory access, 110

ps_pglobal_lookup — look up a symbol in the symbol table of the load object in the target process, 109
 ps_pglobal_sym — look up a symbol in the symbol table of the load object in the target process, 109
 ps_pstop — process and LWP control in libthread_db, 111
 ps_ptread — interfaces in libthread_db that target process memory access, 110
 ps_ptwrite — interfaces in libthread_db that target process memory access, 110
 pthread_atfork — register fork handlers, 113
 pthread_attr_destroy — initialize and destroy threads attribute object, 125
 pthread_attr_getdetachstate — get or set detachstate attribute, 115
 pthread_attr_getguardsize — get or set the thread guardsize attribute, 116
 pthread_attr_getinheritsched — get or set inheritsched attribute, 118
 pthread_attr_getschedparam — get or set schedparam attribute, 120
 pthread_attr_getschedpolicy — get or set schedpolicy attribute, 121
 pthread_attr_getscope — get or set contentionscope attribute, 122
 pthread_attr_getstackaddr — get or set stackaddr attribute, 123
 pthread_attr_getstacksize — get or set stacksize attribute, 124
 pthread_attr_init — initialize and destroy threads attribute object, 125
 pthread_attr_setdetachstate — get or set detachstate attribute, 115
 pthread_attr_setguardsize — get or set the thread guardsize attribute, 116
 pthread_attr_setinheritsched — get or set inheritsched attribute, 118
 pthread_attr_setschedparam — get or set schedparam attribute, 120
 pthread_attr_setschedpolicy — get or set schedpolicy attribute, 121
 pthread_attr_setscope — get or set contentionscope attribute, 122
 pthread_attr_setstackaddr — get or set stackaddr attribute, 123
 pthread_attr_setstacksize — get or set stacksize attribute, 124
 pthread_cleanup_pop — pop a thread cancellation cleanup handler, 128
 pthread_cleanup_push — push a thread cancellation cleanup handler, 129
 pthread_condattr_destroy — initialize or destroy condition variable attributes object, 132
 pthread_condattr_getpshared — get or set the process-shared condition variable attributes, 130
 pthread_condattr_init — initialize or destroy condition variable attributes object, 132
 pthread_condattr_setpshared — get or set the process-shared condition variable attributes, 130
 pthread_cond_broadcast — signal or broadcast a condition, 136
 pthread_cond_destroy — initialize or destroy condition variables, 134
 pthread_cond_init — initialize or destroy condition variables, 134
 pthread_cond_reltimedwait_np — wait on a condition, 138
 pthread_cond_signal — signal or broadcast a condition, 136
 pthread_cond_timedwait — wait on a condition, 138
 pthread_cond_wait — wait on a condition, 138
 pthread_create — create a thread, 141
 pthread_detach — detach a thread, 145
 pthread_equal — compare thread IDs, 146
 pthread_exit — terminate calling thread, 147
 pthread_getconcurrency — get or set level of concurrency, 148
 pthread_getschedparam — access dynamic thread scheduling parameters, 150
 pthread_getspecific — manage thread-specific data, 152
 pthread_join — wait for thread termination, 153
 pthread_key_create — create thread-specific data key, 155
 pthread_key_delete — delete thread-specific data key, 157

`pthread_mutexattr_destroy` — initialize and destroy mutex attributes object, 170
`pthread_mutexattr_getprioceiling` — get and set prioceiling attribute of mutex attribute object, 159
`pthread_mutexattr_getprotocol` — get and set protocol attribute of mutex attribute object, 161
`pthread_mutexattr_getpshared` — get and set process-shared attribute, 164
`pthread_mutexattr_gettype` — get or set a mutex type, 168
`pthread_mutexattr_init` — initialize and destroy mutex attributes object, 170
`pthread_mutexattr_setprioceiling` — get and set prioceiling attribute of mutex attribute object, 159
`pthread_mutexattr_setprotocol` — get and set protocol attribute of mutex attribute object, 161
`pthread_mutexattr_setpshared` — get and set process-shared attribute, 164
`pthread_mutexattr_settype` — get or set a mutex type, 168
`pthread_mutex_consistent_np` — make a mutex consistent after owner death, 171
`pthread_mutex_destroy` — initialize or destroy a mutex, 175
`pthread_mutex_getprioceiling` — change the priority ceiling of a mutex, 173
`pthread_mutex_init` — initialize or destroy a mutex, 175
`pthread_mutex_lock` — lock or unlock a mutex, 177
`pthread_mutex_setprioceiling` — change the priority ceiling of a mutex, 173
`pthread_mutex_trylock` — lock or unlock a mutex, 177
`pthread_mutex_unlock` — lock or unlock a mutex, 177
`pthread_once` — initialize dynamic package, 180
`pthread_rwlockattr_destroy` — initialize and destroy read-write lock attributes object, 182
`pthread_rwlockattr_getpshared` — get and set process-shared attribute of read-write lock attributes object, 181
`pthread_rwlockattr_init` — initialize and destroy read-write lock attributes object, 182
`pthread_rwlockattr_setpshared` — get and set process-shared attribute of read-write lock attributes object, 181
`pthread_rwlock_destroy` — initialize or destroy a read-write lock object, 183
`pthread_rwlock_init` — initialize or destroy a read-write lock object, 183
`pthread_rwlock_rdlock` — lock or attempt to lock a read-write lock object for reading, 185
`pthread_rwlock_tryrdlock` — lock or attempt to lock a read-write lock object for reading, 185
`pthread_rwlock_trywrlock` — lock or attempt to lock a read-write lock object for writing, 188
`pthread_rwlock_unlock` — unlock a read-write lock object, 187
`pthread_rwlock_wrlock` — lock or attempt to lock a read-write lock object for writing, 188

`pthread_self` — get calling thread's ID, 190
`pthread_setcancelstate` — enable or disable cancellation, 191
`pthread_setcanceltype` — set the cancellation type of a thread, 193
`pthread_setconcurrency` — get or set level of concurrency, 148
`pthread_setschedparam` — access dynamic thread scheduling parameters, 150
`pthread_setspecific` — manage thread-specific data, 152
`pthread_sigmask` — change or examine calling thread's signal mask, 195
`pthread_testcancel` — create cancellation point in the calling thread., 200
push a thread cancellation cleanup handler — `pthread_cleanup_push`, 129

R

read or write asynchronous I/O operations
— `aioread`, 23
— `aiowrite`, 23

reading and writing thread registers in `libthread_db`
— `td_thr_getfpregs`, 262
— `td_thr_getgregs`, 262
— `td_thr_getxregs`, 262
— `td_thr_getxregsize`, 262
— `td_thr_setfpregs`, 262
— `td_thr_setgregs`, 262
— `td_thr_setxregs`, 262

register fork handlers — `pthread_atfork`, 113

return credential information associated with the client — `door_cred`, 58

return from a door invocation — `door_return`, 61

return information associated with a door descriptor — `door_info`, 59

return the synchronization handle for the object on which a thread is blocked — `td_thr_sleepinfo`, 270

revoke access to a door descriptor — `door_revoke`, 62

routines that access the target process register in `libthread_db`
— `ps_lgetfpregs`, 107
— `ps_lgetregs`, 107
— `ps_lgetxregs`, 107
— `ps_lgetxregsize`, 107
— `ps_lsetfpregs`, 107
— `ps_lsetregs`, 107
— `ps_lsetxregs`, 107

`rwlock_destroy()` — destroy a readers/writer lock, 201

`rwlock_init()` — initialize a readers/writer lock, 201

`rw_rdlock()` — acquire a read lock, 201

`rw_tryrdlock()` — acquire a read lock, 201

`rw_trywrlock()` — acquire a write lock, 201

`rw_unlock()` — unlock a readers/writer lock, 201

`rw_wrlock()` — acquire a write lock, 201

S

`schedctl_exit` — preemption control, 204

`schedctl_init` — preemption control, 204

`schedctl_lookup` — preemption control, 204

`schedctl_start` — preemption control, 204

`schedctl_stop` — preemption control, 204

`sched_getparam` — get scheduling parameters, 206, 210

`sched_get_priority_max` — get scheduling parameter limits, 207

`sched_get_priority_min` — get scheduling parameter limits, 207

`sched_getscheduler` — get scheduling policy, 208

`sched_rr_get_interval` — get execution time limits, 209

`sched_setparam` — set/get scheduling parameters, 210

`sched_setscheduler` — set scheduling policy and scheduling parameters, 212

`sched_yield` — yield processor, 214

`sema_destroy()` — destroy a semaphore, 215

`sema_init()` — initialize a semaphore, 215

semaphore

- acquire or wait for — `sem_wait`, `sem_trywait`, 230
- close a named one — `sem_close`, 219
- destroy an unnamed one — `sem_destroy`, 220
- get the value — `sem_getvalue`, 221
- increment the count — `sem_post`, 227
- initialize an unnamed one — `sem_init`, 222
- initialize/open a named one — `sem_open`, 224
- remove a named one — `sem_unlink`, 229

`sema_post()` — increment a semaphore, 215

`sema_trywait()` — decrement a semaphore, 215

`sema_wait()` — decrement a semaphore, 215

`sem_close` — close a named semaphore, 219

`sem_destroy` — destroy an unnamed semaphore, 220

`sem_getvalue` — get the value of a semaphore, 221

`sem_init` — initialize an unnamed semaphore, 222

`sem_open` — initialize/open a named semaphore, 224

- sem_post — increment the count of a semaphore, 227
- sem_trywait — acquire or wait for a semaphore, 230
- sem_unlink — remove a named semaphore, 229
- sem_wait — acquire or wait for a semaphore, 230
- set concurrency level for target process — td_ta_setconcurrency, 258
- set scheduling policy and scheduling parameters — sched_setscheduler, 212
- set the cancellation type of a thread — pthread_setcanceltype, 193
- set the priority of a thread — td_thr_setprio, 268
- set/get scheduling parameters
 - sched_getparam, 210
 - sched_setparam, 210
- shared memory object
 - open — shm_open, 233
 - remove — shm_unlink, 236
- shm_open — open a shared memory object, 233
- shm_unlink — remove a shared memory object, 236
- signal
 - queue one to a process — sigqueue, 237
 - wait for queued signals — sigwaitinfo, sigtimedwait, 239
- signal or broadcast a condition
 - pthread_cond_broadcast, 136
 - pthread_cond_signal, 136
- sigqueue — queue a signal to a process, 237
- sigtimedwait — wait for queued signals, 239
- sigwaitinfo — wait for queued signals, 239
- sleep
 - high resolution — nanosleep, 103
- specify an alternative door server thread creation function — door_server_create, 63
- suspend and resume threads in libthread_db
 - td_thr_dbresume, 261
 - td_thr_dbsuspend, 261
- synchronize a file's data
 - fdatsync, 65

T

- tda_ta_clear_event — thread events in libthread_db, 249
- td_event_addset — thread events in libthread_db, 249
- td_event_delset — thread events in libthread_db, 249
- td_event_emptyset — thread events in libthread_db, 249
- td_event_fillset — thread events in libthread_db, 249
- td_eventisempty — thread events in libthread_db, 249
- td_eventismember — thread events in libthread_db, 249
- td_init — initialization function for libthread_db library of interfaces, 241
- td_log — placeholder for future logging functionality, 242
- td_sync_get_info — operations on a synchronization object in libthread_db, 243
- td_sync_get_stats — operations on a synchronization object in libthread_db, 243
- td_sync_setstate — operations on a synchronization object in libthread_db, 243
- td_sync_waiters — operations on a synchronization object in libthread_db, 243
- td_ta_delete — allocate and deallocate process handles for libthread_db, 256
- td_ta_enable_stats — collect target process statistics for libthread_db, 247
- td_ta_event_addr — thread events in libthread_db, 249
- td_ta_event_getmsg — thread events in libthread_db, 249
- td_ta_event_getmsg — thread events in libthread_db, 249
- td_ta_get_nthreads — gets the total number of threads in a process for libthread_db, 253
- td_ta_get_ph — allocate and deallocate process handles for libthread_db, 256
- td_ta_get_stats — collect target process statistics for libthread_db, 247
- td_ta_map_addr2sync — get a synchronization object handle from a synchronization object's address, 254
- td_ta_map_addr2thr — convert a thread id or thread address to a thread handle, 255

td_ta_map_id2thr — convert a thread id or thread address to a thread handle, 255
 td_ta_new — allocate and deallocate process handles for libthread_db, 256
 td_ta_reset_stats — collect target process statistics for libthread_db, 247
 td_ta_setconcurrency — set concurrency level for target process, 258
 td_ta_set_event — thread events in libthread_db, 249
 td_ta_sync_iter — iterator functions on process handles from libthread_db, 259
 td_ta_sync_tracking_enable — operations on a synchronization object in libthread_db, 243
 td_ta_thr_iter — iterator functions on process handles from libthread_db, 259
 td_ta_tsd_iter — iterator functions on process handles from libthread_db, 259
 td_thr_clear_event — thread events in libthread_db, 249
 td_thr_dbresume — suspend and resume threads in libthread_db, 261
 td_thr_dbsuspend — suspend and resume threads in libthread_db, 261
 td_thr_event_enable — thread events in libthread_db, 249
 td_thr_event_getmsg — thread events in libthread_db, 249
 td_thr_getfpregs — reading and writing thread registers in libthread_db, 262
 td_thr_getgregs — reading and writing thread registers in libthread_db, 262
 Intel IA, 263
 SPARC, 263
 td_thr_get_info — get thread information in libthread_db library of interfaces, 264
 td_thr_getxregs — reading and writing thread registers in libthread_db, 262
 td_thr_getxregsize — reading and writing thread registers in libthread_db, 262
 td_thr_lockowner — iterate over the set of locks owned by a thread, 267
 td_thr_set_event — thread events in libthread_db, 249
 td_thr_setfpregs — reading and writing thread registers in libthread_db, 262
 td_thr_setgregs — reading and writing thread registers in libthread_db, 262
 td_thr_setprio — set the priority of a thread, 268
 td_thr_setsigpending — manage thread signals for libthread_db, 269
 td_thr_setxregs — reading and writing thread registers in libthread_db, 262
 td_thr_sigsetmask — manage thread signals for libthread_db, 269
 td_thr_sleepinfo — return the synchronization handle for the object on which a thread is blocked, 270
 td_thr_tsd — get a thread's thread-specific data for libthread_db library of interfaces, 271
 td_thr_validate — test a thread handle for validity, 272
 terminate calling thread — pthread_exit, 147
 terminate the calling thread — thr_exit, 286
 test a thread handle for validity —
 td_thr_validate, 272
 thr_continue — continue thread execution, 305
 thr_create — create a thread, 273
 thread events in libthread_db
 — tda_ta_clear_event, 249
 — td_event_addset, 249
 — td_event_delset, 249
 — td_event_emptyset, 249
 — td_event_fillset, 249
 — td_eventisempty, 249
 — td_eventismember, 249
 — td_ta_event_addr, 249
 — td_ta_event_getmsg, 249
 — td_ta_set_event, 249
 — td_thr_clear_event, 249
 — td_thr_event_enable, 249
 — td_thr_event_getmsg, 249
 — td_thr_set_event, 249
 thread yield to another thread —
 thr_yield, 306
 threads — concepts related to POSIX pthreads and Solaris threads and the libpthread and libthread libraries, 279
 thread-specific-data functions
 — thr_getspecific, 292
 — thr_keycreate, 292
 — thr_setspecific, 292

`thr_exit` — terminate the calling thread, 286
`thr_getconcurrency` — get thread concurrency level, 288
`thr_getprio` — access dynamic thread scheduling, 289
`thr_getspecific` — thread-specific-data functions, 292
`thr_join` — wait for thread termination, 290
`thr_keycreate` — thread-specific-data functions, 292
 Create Key, 292
 Get Value, 292
 Set Value, 292
`thr_main` — identifies the calling thread as the main thread or not the main thread, 296
`thr_self` — get calling thread's ID, 298
`thr_setconcurrency` — set thread concurrency level, 288
`thr_setprio` — access dynamic thread scheduling, 289
`thr_setspecific` — thread-specific-data functions, 292
`thr_sigsetmask` — change or examine calling thread's signal mask, 299
`thr_stksegment` — get thread stack address and size, 304
`thr_suspend` — suspend thread execution, 305
`thr_yield` — thread yield to another thread, 306
`timer_getoverrun` — per-process timers, 310
`timer_gettime` — per-process timers, 310
`timer_settime` — per-process timers, 310

U

unlock a read-write lock object —
 `pthread_rwlock_unlock`, 187

W

wait on a condition — `pthread_cond_wait`, 138
wait for thread termination —
 `pthread_join`, 153, 290

Y

yield processor — `sched_yield`, 214

