

17643 TSP defined for SCHED_FIFO, SCHED_RR, and SCHED_SPORADIC shall be used in an
17644 implementation-defined manner. Each thread with system scheduling contention scope
17645 competes for the processors in its scheduling allocation domain in an implementation-defined
17646 manner according to its priority. Threads with process scheduling contention scope are
17647 scheduled relative to other threads within the same scheduling contention scope in the process.
17648 TSP If _POSIX_THREAD_SPORADIC_SERVER is defined, the rules defined for SCHED_SPORADIC
17649 in Scheduling Policies (on page 501) shall be used in an implementation-defined manner for
17650 application threads whose scheduling allocation domain size is greater than one.

17651 Scheduling Documentation

17652 If _POSIX_PRIORITY_SCHEDULING is defined, then any scheduling policies beyond
17653 TSP SCHED_OTHER, SCHED_FIFO, SCHED_RR, and SCHED_SPORADIC, as well as the effects of
17654 the scheduling policies indicated by these other values, and the attributes required in order to
17655 support such a policy, are implementation-defined. Furthermore, the implementation shall
17656 document the effect of all processor scheduling allocation domain values supported for these
17657 policies.

17658 2.9.5 Thread Cancellation

17659 The thread cancellation mechanism allows a thread to terminate the execution of any other
17660 thread in the process in a controlled manner. The target thread (that is, the one that is being
17661 canceled) is allowed to hold cancellation requests pending in a number of ways and to perform
17662 application-specific cleanup processing when the notice of cancellation is acted upon.

17663 Cancellation is controlled by the cancellation control functions. Each thread maintains its own
17664 cancelability state. Cancellation may only occur at cancellation points or when the thread is
17665 asynchronously cancelable.

17666 The thread cancellation mechanism described in this section depends upon programs having set
17667 *deferred* cancelability state, which is specified as the default. Applications shall also carefully
17668 follow static lexical scoping rules in their execution behavior. For example, use of *setjmp()*,
17669 *return*, *goto*, and so on, to leave user-defined cancellation scopes without doing the necessary
17670 scope pop operation results in undefined behavior.

17671 Use of asynchronous cancelability while holding resources which potentially need to be released
17672 may result in resource loss. Similarly, cancellation scopes may only be safely manipulated
17673 (pushed and popped) when the thread is in the *deferred* or *disabled* cancelability states.

17674 2.9.5.1 Cancelability States

17675 The cancelability state of a thread determines the action taken upon receipt of a cancellation
17676 request. The thread may control cancellation in a number of ways.

17677 Each thread maintains its own cancelability state, which may be encoded in two bits:

- 17678 1. Cancelability-Enable: When cancelability is PTHREAD_CANCEL_DISABLE (as defined
17679 in the Base Definitions volume of POSIX.1-2008, *<pthread.h>*), cancellation requests
17680 against the target thread are held pending. By default, cancelability is set to
17681 PTHREAD_CANCEL_ENABLE (as defined in *<pthread.h>*).
- 17682 2. Cancelability Type: When cancelability is enabled and the cancelability type is
17683 PTHREAD_CANCEL_ASYNCHRONOUS (as defined in *<pthread.h>*), new or pending
17684 cancellation requests may be acted upon at any time. When cancelability is enabled and
17685 the cancelability type is PTHREAD_CANCEL_DEFERRED (as defined in *<pthread.h>*),

17686 cancellation requests are held pending until a cancellation point (see below) is reached. If
 17687 cancelability is disabled, the setting of the cancelability type has no immediate effect as all
 17688 cancellation requests are held pending; however, once cancelability is enabled again the
 17689 new type is in effect. The cancelability type is PTHREAD_CANCEL_DEFERRED in all
 17690 newly created threads including the thread in which *main()* was first invoked.

17691 2.9.5.2 *Cancellation Points*

17692 Cancellation points shall occur when a thread is executing the following functions:

17693	<i>accept()</i>	<i>nanosleep()</i>	<i>select()</i>
17694	<i>aio_suspend()</i>	<i>open()</i>	<i>sem_timedwait()</i>
17695	<i>clock_nanosleep()</i>	<i>openat()</i>	<i>sem_wait()</i>
17696	<i>close()</i>	<i>pause()</i>	<i>send()</i>
17697	<i>connect()</i>	<i>poll()</i>	<i>sendmsg()</i>
17698	<i>creat()</i>	<i>pread()</i>	<i>sendto()</i>
17699	<i>fcntl()</i> †	<i>pselect()</i>	<i>sigsuspend()</i>
17700	<i>fdatasync()</i>	<i>pthread_cond_timedwait()</i>	<i>sigtimedwait()</i>
17701	<i>fsync()</i>	<i>pthread_cond_wait()</i>	<i>sigwait()</i>
17702	<i>getmsg()</i>	<i>pthread_join()</i>	<i>sigwaitinfo()</i>
17703	<i>getpmsg()</i>	<i>pthread_testcancel()</i>	<i>sleep()</i>
17704	<i>lockf()</i> ††	<i>putmsg()</i>	<i>system()</i>
17705	<i>mq_receive()</i>	<i>putpmsg()</i>	<i>tcdrain()</i>
17706	<i>mq_send()</i>	<i>pwrite()</i>	<i>wait()</i>
17707	<i>mq_timedreceive()</i>	<i>read()</i>	<i>waitid()</i>
17708	<i>mq_timedsend()</i>	<i>readv()</i>	<i>waitpid()</i>
17709	<i>msgrcv()</i>	<i>recv()</i>	<i>write()</i>
17710	<i>msgsnd()</i>	<i>recvfrom()</i>	<i>writev()</i>
17711	<i>msync()</i>	<i>recvmsg()</i>	

17712 † When the *cmd* argument is F_SETLKW.

17713 †† When the *function* argument is F_LOCK.

17714 A cancellation point may also occur when a thread is executing the following functions:

17715	access()	fprintf()	getprotobyname()
17716	asctime()	fputc()	getprotoent()
17717	asctime_r()	fputs()	getpwent()
17718	catclose()	fputwc()	getpwnam()
17719	catgets()	fputws()	getpwnam_r()
17720	catopen()	fread()	getpwuid()
17721	chmod()	freopen()	getpwuid_r()
17722	chown()	fscanf()	gets()
17723	closedir()	fseek()	getservbyname()
17724	closelog()	fseeko()	getservbyport()
17725	ctermid()	fsetpos()	getservent()
17726	ctime()	fstat()	getutxent()
17727	ctime_r()	fstatat()	getutxid()
17728	dbm_close()	ftell()	getutxline()
17729	dbm_delete()	ftello()	getwc()
17730	dbm_fetch()	ftw()	getwchar()
17731	dbm_nextkey()	futimens()	glob()
17732	dbm_open()	fwprintf()	iconv_close()
17733	dbm_store()	fwrite()	iconv_open()
17734	dlclose()	fwscanf()	ioctl()
17735	dlopen()	getaddrinfo()	link()
17736	dprintf()	getc()	linkat()
17737	endgrent()	getc_unlocked()	lio_listio()
17738	endhostent()	getchar()	localtime()
17739	endnetent()	getchar_unlocked()	localtime_r()
17740	endprotoent()	getcwd()	lockf()
17741	endpwent()	getdate()	lseek()
17742	endservent()	getdelim()	lstat()
17743	endutxent()	getgrent()	mkdir()
17744	faccessat()	getgrgid()	mkdirat()
17745	fchmod()	getgrgid_r()	mkdtemp()
17746	fchmodat()	getgrnam()	mkfifo()
17747	fchown()	getgrnam_r()	mkfifoat()
17748	fchownat()	gethostent()	mknod()
17749	fclose()	gethostid()	mknodat()
17750	fcntl()†	gethostname()	mkstemp()
17751	fflush()	getline()	mktimes()
17752	fgetc()	getlogin()	nftw()
17753	fgetpos()	getlogin_r()	opendir()
17754	fgets()	getnameinfo()	openlog()
17755	fgetwc()	getnetbyaddr()	pathconf()
17756	fgetws()	getnetbyname()	pclose()
17757	fmtmsg()	getnetent()	perror()
17758	fopen()	getopt()††	popen()
17759	fpathconf()	getprotobyname()	posix_fadvise()

17760	<i>posix_fallocate()</i>	<i>putc()</i>	<i>strerror()</i>
17761	<i>posix_madvise()</i>	<i>putc_unlocked()</i>	<i>strerror_r()</i>
17762	<i>posix_openpt()</i>	<i>putchar()</i>	<i>strftime()</i>
17763	<i>posix_spawn()</i>	<i>putchar_unlocked()</i>	<i>symlink()</i>
17764	<i>posix_spawnnp()</i>	<i>puts()</i>	<i>symlinkat()</i>
17765	<i>posix_trace_clear()</i>	<i>pututxline()</i>	<i>sync()</i>
17766	<i>posix_trace_close()</i>	<i>putwc()</i>	<i>syslog()</i>
17767	<i>posix_trace_create()</i>	<i>putwchar()</i>	<i>tmpfile()</i>
17768	<i>posix_trace_create_withlog()</i>	<i>readdir()</i>	<i>tmpnam()</i>
17769	<i>posix_trace_eventtypelist_getnext_id()</i>	<i>readdir_r()</i>	<i>ttyname()</i>
17770	<i>posix_trace_eventtypelist_rewind()</i>	<i>readlink()</i>	<i>ttyname_r()</i>
17771	<i>posix_trace_flush()</i>	<i>readlinkat()</i>	<i>tzset()</i>
17772	<i>posix_trace_get_attr()</i>	<i>remove()</i>	<i>ungetc()</i>
17773	<i>posix_trace_get_filter()</i>	<i>rename()</i>	<i>ungetwc()</i>
17774	<i>posix_trace_get_status()</i>	<i>renameat()</i>	<i>unlink()</i>
17775	<i>posix_trace_getnext_event()</i>	<i>rewind()</i>	<i>unlinkat()</i>
17776	<i>posix_trace_open()</i>	<i>rewinddir()</i>	<i>utime()</i>
17777	<i>posix_trace_rewind()</i>	<i>scandir()</i>	<i>utimensat()</i>
17778	<i>posix_trace_set_filter()</i>	<i>scanf()</i>	<i>utimes()</i>
17779	<i>posix_trace_shutdown()</i>	<i>seekdir()</i>	<i>vdprintf()</i>
17780	<i>posix_trace_timedgetnext_event()</i>	<i>semop()</i>	<i>vfprintf()</i>
17781	<i>posix_typed_mem_open()</i>	<i>setgrent()</i>	<i>vfwprintf()</i>
17782	<i>printf()</i>	<i>sethostent()</i>	<i>vprintf()</i>
17783	<i>psiginfo()</i>	<i>setnetent()</i>	<i>vwprintf()</i>
17784	<i>psignal()</i>	<i>setprotoent()</i>	<i>wcsftime()</i>
17785	<i>pthread_rwlock_rdlock()</i>	<i>setpwent()</i>	<i>wordexp()</i>
17786	<i>pthread_rwlock_timedrdlock()</i>	<i>setservent()</i>	<i>wprintf()</i>
17787	<i>pthread_rwlock_timedwrlock()</i>	<i>setutxent()</i>	<i>wscanf()</i>
17788	<i>pthread_rwlock_wrlock()</i>	<i>sigpause()</i>	
17789		<i>stat()</i>	

17790 An implementation shall not introduce cancellation points into any other functions specified in
 17791 this volume of POSIX.1-2008.

17792 The side-effects of acting upon a cancellation request while suspended during a call of a function
 17793 are the same as the side-effects that may be seen in a single-threaded program when a call to a
 17794 function is interrupted by a signal and the given function returns [EINTR]. Any such side-
 17795 effects occur before any cancellation cleanup handlers are called.

17796 Whenever a thread has cancelability enabled and a cancellation request has been made with that
 17797 thread as the target, and the thread then calls any function that is a cancellation point (such as
 17798 *pthread_testcancel()* or *read()*), the cancellation request shall be acted upon before the function
 17799 returns. If a thread has cancelability enabled and a cancellation request is made with the thread
 17800 as a target while the thread is suspended at a cancellation point, the thread shall be awakened
 17801 and the cancellation request shall be acted upon. It is unspecified whether the cancellation
 17802 request is acted upon or whether the cancellation request remains pending and the thread
 17803 resumes normal execution if:

17804 • The thread is suspended at a cancellation point and the event for which it is waiting occurs

17805 † For any value of the *cmd* argument.

17806 †† If *opterr* is non-zero.

17807 • A specified timeout expired
17808 before the cancellation request is acted upon.

17809 2.9.5.3 *Thread Cancellation Cleanup Handlers*

17810 Each thread maintains a list of cancellation cleanup handlers. The programmer uses the
17811 *pthread_cleanup_push()* and *pthread_cleanup_pop()* functions to place routines on and remove
17812 routines from this list.

17813 When a cancellation request is acted upon, or when a thread calls *pthread_exit()*, the thread first
17814 disables cancellation by setting its cancelability state to PTHREAD_CANCEL_DISABLE and its
17815 cancelability type to PTHREAD_CANCEL_DEFERRED. The cancelability state shall remain set to
17816 PTHREAD_CANCEL_DISABLE until the thread has terminated. The behavior is undefined if
17817 a cancellation cleanup handler or thread-specific data destructor routine changes the
17818 cancelability state to PTHREAD_CANCEL_ENABLE.

17819 The routines in the thread's list of cancellation cleanup handlers are invoked one by one in LIFO
17820 sequence; that is, the last routine pushed onto the list (Last In) is the first to be invoked (First
17821 Out). When the cancellation cleanup handler for a scope is invoked, the storage for that scope
17822 remains valid. If the last cancellation cleanup handler returns, thread-specific data destructors (if
17823 any) associated with thread-specific data keys for which the thread has non-NULL values will
17824 be run, in unspecified order, as described for *pthread_key_create()*.

17825 After all cancellation cleanup handlers and thread-specific data destructors have returned,
17826 thread execution is terminated. If the thread has terminated because of a call to *pthread_exit()*,
17827 the *value_ptr* argument is made available to any threads joining with the target. If the thread has
17828 terminated by acting on a cancellation request, a status of PTHREAD_CANCELED is made
17829 available to any threads joining with the target. The symbolic constant PTHREAD_CANCELED
17830 expands to a constant expression of type (**void** *) whose value matches no pointer to an object in
17831 memory nor the value NULL.

17832 A side-effect of acting upon a cancellation request while in a condition variable wait is that the
17833 mutex is re-acquired before calling the first cancellation cleanup handler. In addition, the thread
17834 is no longer considered to be waiting for the condition and the thread shall not have consumed
17835 any pending condition signals on the condition.

17836 A cancellation cleanup handler cannot exit via *longjmp()* or *siglongjmp()*.

17837 2.9.5.4 *Async-Cancel Safety*

17838 The *pthread_cancel()*, *pthread_setcancelstate()*, and *pthread_setcanceltype()* functions are defined to
17839 be async-cancel safe.

17840 No other functions in this volume of POSIX.1-2008 are required to be async-cancel-safe.

17841 2.9.6 **Thread Read-Write Locks**

17842 Multiple readers, single writer (read-write) locks allow many threads to have simultaneous
17843 read-only access to data while allowing only one thread to have exclusive write access at any
17844 given time. They are typically used to protect data that is read more frequently than it is
17845 changed.

17846 One or more readers acquire read access to the resource by performing a read lock operation on
17847 the associated read-write lock. A writer acquires exclusive write access by performing a write
17848 lock operation. Basically, all readers exclude any writers and a writer excludes all readers and