

**NAME**

**taskqueue** - asynchronous task execution

**SYNOPSIS**

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

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

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

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

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

```
typedef void (*task_fn_t)(void *context, int pending);
```

```
typedef void (*taskqueue_enqueue_fn)(void *context);
```

```
struct task {
    STAILQ_ENTRY(task)    ta_link; /* link for queue */
    u_short              ta_pending; /* count times queued */
    u_short              ta_priority; /* priority of task in queue */
    task_fn_t            ta_func; /* task handler */
    void                 *ta_context; /* argument for handler */
};
```

```
enum taskqueue_callback_type {
    TASKQUEUE_CALLBACK_TYPE_INIT,
    TASKQUEUE_CALLBACK_TYPE_SHUTDOWN,
};
```

```
typedef void (*taskqueue_callback_fn)(void *context);
```

```
struct timeout_task;
```

```
struct taskqueue *
```

```
taskqueue_create(const char *name, int mflags, taskqueue_enqueue_fn enqueue, void *context);
```

```
struct taskqueue *
```

```
taskqueue_create_fast(const char *name, int mflags, taskqueue_enqueue_fn enqueue, void *context);
```

```
int
```

```
taskqueue_start_threads(struct taskqueue **tqp, int count, int pri, const char *name, ...);
```

```
int
```

```
taskqueue_start_threads_cpuset(struct taskqueue **tqp, int count, int pri, cpuset_t *mask,  
const char *name, ...);
```

*int*

```
taskqueue_start_threads_in_proc(struct taskqueue **tqp, int count, int pri, struct proc *proc,  
const char *name, ...);
```

*void*

```
taskqueue_set_callback(struct taskqueue *queue, enum taskqueue_callback_type cb_type,  
taskqueue_callback_fn callback, void *context);
```

*void*

```
taskqueue_free(struct taskqueue *queue);
```

*int*

```
taskqueue_enqueue(struct taskqueue *queue, struct task *task);
```

*int*

```
taskqueue_enqueue_flags(struct taskqueue *queue, struct task *task, int flags);
```

*int*

```
taskqueue_enqueue_timeout(struct taskqueue *queue, struct timeout_task *timeout_task, int ticks);
```

*int*

```
taskqueue_enqueue_timeout_sbt(struct taskqueue *queue, struct timeout_task *timeout_task,  
sbintime_t sbt, sbintime_t pr, int flags);
```

*int*

```
taskqueue_cancel(struct taskqueue *queue, struct task *task, u_int *pendp);
```

*int*

```
taskqueue_cancel_timeout(struct taskqueue *queue, struct timeout_task *timeout_task, u_int *pendp);
```

*void*

```
taskqueue_drain(struct taskqueue *queue, struct task *task);
```

*void*

```
taskqueue_drain_timeout(struct taskqueue *queue, struct timeout_task *timeout_task);
```

*void*

```
taskqueue_drain_all(struct taskqueue *queue);
```

*void*

```
taskqueue_quiesce(struct taskqueue *queue);
```

*void*

```
taskqueue_block(struct taskqueue *queue);
```

*void*

```
taskqueue_unblock(struct taskqueue *queue);
```

*int*

```
taskqueue_member(struct taskqueue *queue, struct thread *td);
```

*void*

```
taskqueue_run(struct taskqueue *queue);
```

```
TASK_INIT(struct task *task, int priority, task_fn_t func, void *context);
```

```
TASK_INITIALIZER(int priority, task_fn_t func, void *context);
```

```
TASKQUEUE_DECLARE(name);
```

```
TASKQUEUE_DEFINE(name, taskqueue_enqueue_fn enqueue, void *context, init);
```

```
TASKQUEUE_FAST_DEFINE(name, taskqueue_enqueue_fn enqueue, void *context, init);
```

```
TASKQUEUE_DEFINE_THREAD(name);
```

```
TASKQUEUE_FAST_DEFINE_THREAD(name);
```

```
TIMEOUT_TASK_INIT(struct taskqueue *queue, struct timeout_task *timeout_task, int priority,  
task_fn_t func, void *context);
```

## DESCRIPTION

These functions provide a simple interface for asynchronous execution of code.

The function **taskqueue\_create()** is used to create new queues. The arguments to **taskqueue\_create()** include a name that should be unique, a set of **malloc(9)** flags that specify whether the call to **malloc()** is allowed to sleep, a function that is called from **taskqueue\_enqueue()** when a task is added to the queue,

and a pointer to the memory location where the identity of the thread that services the queue is recorded. The function called from **taskqueue\_enqueue()** must arrange for the queue to be processed (for instance by scheduling a software interrupt or waking a kernel thread). The memory location where the thread identity is recorded is used to signal the service thread(s) to terminate--when this value is set to zero and the thread is signaled it will terminate. If the queue is intended for use in fast interrupt handlers **taskqueue\_create\_fast()** should be used in place of **taskqueue\_create()**.

The function **taskqueue\_free()** should be used to free the memory used by the queue. Any tasks that are on the queue will be executed at this time after which the thread servicing the queue will be signaled that it should exit.

Once a taskqueue has been created, its threads should be started using **taskqueue\_start\_threads()**, **taskqueue\_start\_threads\_cpuset()** or **taskqueue\_start\_threads\_in\_proc()**. **taskqueue\_start\_threads\_cpuset()** takes a *cpuset* argument which will cause the threads which are started for the taskqueue to be restricted to run on the given CPUs. **taskqueue\_start\_threads\_in\_proc()** takes a *proc* argument which will cause the threads which are started for the taskqueue to be assigned to the given kernel process. Callbacks may optionally be registered using **taskqueue\_set\_callback()**. Currently, callbacks may be registered for the following purposes:

TASKQUEUE_CALLBACK_TYPE_INIT	This callback is called by every thread in the taskqueue, before it executes any tasks. This callback must be set before the taskqueue's threads are started.
TASKQUEUE_CALLBACK_TYPE_SHUTDOWN	This callback is called by every thread in the taskqueue, after it executes its last task. This callback will always be called before the taskqueue structure is reclaimed.

To add a task to the list of tasks queued on a taskqueue, call **taskqueue\_enqueue()** with pointers to the queue and task. If the task's *ta\_pending* field is non-zero, then it is simply incremented to reflect the number of times the task was enqueued, up to a cap of USHRT\_MAX. Otherwise, the task is added to the list before the first task which has a lower *ta\_priority* value or at the end of the list if no tasks have a lower priority. Enqueueing a task does not perform any memory allocation which makes it suitable for calling from an interrupt handler. This function will return EPIPE if the queue is being freed.

When a task is executed, first it is removed from the queue, the value of *ta\_pending* is recorded and then the field is zeroed. The function *ta\_func* from the task structure is called with the value of the field *ta\_context* as its first argument and the value of *ta\_pending* as its second argument. After the function *ta\_func* returns, **wakeup(9)** is called on the task pointer passed to **taskqueue\_enqueue()**.

The **taskqueue\_enqueue\_flags()** accepts an extra *flags* parameter which specifies a set of optional flags to alter the behavior of **taskqueue\_enqueue()**. It contains one or more of the following flags:

**TASKQUEUE\_FAIL\_IF\_PENDING**      **taskqueue\_enqueue\_flags()** fails if the task is already scheduled for execution. **EEXIST** is returned and the *ta\_pending* counter value remains unchanged.

**TASKQUEUE\_FAIL\_IF\_CANCELING** **taskqueue\_enqueue\_flags()** fails if the task is in the canceling state and **ECANCELED** is returned.

The **taskqueue\_enqueue\_timeout()** function is used to schedule the enqueue after the specified number of *ticks*. The **taskqueue\_enqueue\_timeout\_sbt()** function provides finer control over the scheduling based on *sbt*, *pr*, and *flags*, as detailed in `callout(9)`. If the *ticks* argument is negative, the already scheduled enqueueing is not re-scheduled. Otherwise, the task is scheduled for enqueueing in the future, after the absolute value of *ticks* is passed. This function returns -1 if the task is being drained. Otherwise, the number of pending calls is returned.

The **taskqueue\_cancel()** function is used to cancel a task. The *ta\_pending* count is cleared, and the old value returned in the reference parameter *pendp*, if it is non-NULL. If the task is currently running, **EBUSY** is returned, otherwise 0. To implement a blocking **taskqueue\_cancel()** that waits for a running task to finish, it could look like:

```
while (taskqueue_cancel(tq, task, NULL) != 0)
    taskqueue_drain(tq, task);
```

Note that, as with **taskqueue\_drain()**, the caller is responsible for ensuring that the task is not re-enqueued after being canceled.

Similarly, the **taskqueue\_cancel\_timeout()** function is used to cancel the scheduled task execution.

The **taskqueue\_drain()** function is used to wait for the task to finish, and the **taskqueue\_drain\_timeout()** function is used to wait for the scheduled task to finish. There is no guarantee that the task will not be enqueued after call to **taskqueue\_drain()**. If the caller wants to put the task into a known state, then before calling **taskqueue\_drain()** the caller should use out-of-band means to ensure that the task would not be enqueued. For example, if the task is enqueued by an interrupt filter, then the interrupt could be disabled.

The **taskqueue\_drain\_all()** function is used to wait for all pending and running tasks that are enqueued on the taskqueue to finish. Tasks posted to the taskqueue after **taskqueue\_drain\_all()** begins processing, including pending enqueues scheduled by a previous call to **taskqueue\_enqueue\_timeout()**, do not

extend the wait time of **taskqueue\_drain\_all()** and may complete after **taskqueue\_drain\_all()** returns. The **taskqueue\_quiesce()** function is used to wait for the queue to become empty and for all running tasks to finish. To avoid blocking indefinitely, the caller must ensure by some mechanism that tasks will eventually stop being posted to the queue.

The **taskqueue\_block()** function blocks the taskqueue. It prevents any enqueued but not running tasks from being executed. Future calls to **taskqueue\_enqueue()** will enqueue tasks, but the tasks will not be run until **taskqueue\_unblock()** is called. Please note that **taskqueue\_block()** does not wait for any currently running tasks to finish. Thus, the **taskqueue\_block()** does not provide a guarantee that **taskqueue\_run()** is not running after **taskqueue\_block()** returns, but it does provide a guarantee that **taskqueue\_run()** will not be called again until **taskqueue\_unblock()** is called. If the caller requires a guarantee that **taskqueue\_run()** is not running, then this must be arranged by the caller. Note that if **taskqueue\_drain()** is called on a task that is enqueued on a taskqueue that is blocked by **taskqueue\_block()**, then **taskqueue\_drain()** can not return until the taskqueue is unblocked. This can result in a deadlock if the thread blocked in **taskqueue\_drain()** is the thread that is supposed to call **taskqueue\_unblock()**. Thus, use of **taskqueue\_drain()** after **taskqueue\_block()** is discouraged, because the state of the task can not be known in advance. The same caveat applies to **taskqueue\_drain\_all()**.

The **taskqueue\_unblock()** function unblocks the previously blocked taskqueue. All enqueued tasks can be run after this call.

The **taskqueue\_member()** function returns 1 if the given thread *td* is part of the given taskqueue *queue* and 0 otherwise.

The **taskqueue\_run()** function will run all pending tasks in the specified *queue*. Normally this function is only used internally.

A convenience macro, **TASK\_INIT(task, priority, func, context)** is provided to initialise a *task* structure. The **TASK\_INITIALIZER()** macro generates an initializer for a task structure. A macro **TIMEOUT\_TASK\_INIT(queue, timeout\_task, priority, func, context)** initializes the *timeout\_task* structure. The values of *priority*, *func*, and *context* are simply copied into the task structure fields and the *ta\_pending* field is cleared.

Five macros **TASKQUEUE\_DECLARE(name)**, **TASKQUEUE\_DEFINE(name, enqueue, context, init)**, **TASKQUEUE\_FAST\_DEFINE(name, enqueue, context, init)**, and **TASKQUEUE\_DEFINE\_THREAD(name)** **TASKQUEUE\_FAST\_DEFINE\_THREAD(name)** are used to declare a reference to a global queue, to define the implementation of the queue, and declare a queue that uses its own thread. The **TASKQUEUE\_DEFINE()** macro arranges to call **taskqueue\_create()** with the values of its *name*, *enqueue* and *context* arguments during system initialisation. After calling **taskqueue\_create()**, the *init* argument to the macro is executed as a C statement, allowing any further

initialisation to be performed (such as registering an interrupt handler, etc.).

The **TASKQUEUE\_DEFINE\_THREAD()** macro defines a new taskqueue with its own kernel thread to serve tasks. The variable *struct taskqueue \*taskqueue\_name* is used to enqueue tasks onto the queue.

**TASKQUEUE\_FAST\_DEFINE()** and **TASKQUEUE\_FAST\_DEFINE\_THREAD()** act just like **TASKQUEUE\_DEFINE()** and **TASKQUEUE\_DEFINE\_THREAD()** respectively but taskqueue is created with **taskqueue\_create\_fast()**.

### Predefined Task Queues

The system provides four global taskqueues, *taskqueue\_fast*, *taskqueue\_swi*, *taskqueue\_swi\_giant*, and *taskqueue\_thread*. The *taskqueue\_fast* queue is for swi handlers dispatched from fast interrupt handlers, where sleep mutexes cannot be used. The swi taskqueues are run via a software interrupt mechanism. The *taskqueue\_swi* queue runs without the protection of the *Giant* kernel lock, and the *taskqueue\_swi\_giant* queue runs with the protection of the *Giant* kernel lock. The thread taskqueue *taskqueue\_thread* runs in a kernel thread context, and tasks run from this thread do not run under the *Giant* kernel lock. If the caller wants to run under *Giant*, he should explicitly acquire and release *Giant* in his taskqueue handler routine.

To use these queues, call **taskqueue\_enqueue()** with the value of the global taskqueue variable for the queue you wish to use.

The software interrupt queues can be used, for instance, for implementing interrupt handlers which must perform a significant amount of processing in the handler. The hardware interrupt handler would perform minimal processing of the interrupt and then enqueue a task to finish the work. This reduces to a minimum the amount of time spent with interrupts disabled.

The thread queue can be used, for instance, by interrupt level routines that need to call kernel functions that do things that can only be done from a thread context. (e.g., call malloc with the M\_WAITOK flag.)

Note that tasks queued on shared taskqueues such as *taskqueue\_swi* may be delayed an indeterminate amount of time before execution. If queueing delays cannot be tolerated then a private taskqueue should be created with a dedicated processing thread.

### SEE ALSO

callout(9), ithread(9), kthread(9), swi(9)

### HISTORY

This interface first appeared in FreeBSD 5.0. There is a similar facility called *work\_queue* in the Linux

kernel.

**AUTHORS**

This manual page was written by Doug Rabson.