

**NAME**

**rmlock**, **rm\_init**, **rm\_init\_flags**, **rm\_destroy**, **rm\_rlock**, **rm\_try\_rlock**, **rm\_wlock**, **rm\_runlock**, **rm\_wunlock**, **rm\_wowned**, **rm\_sleep**, **rm\_assert**, **RM\_SYSINIT**, **RM\_SYSINIT\_FLAGS**, **rms\_init**, **rms\_destroy**, **rms\_rlock**, **rms\_wlock**, **rms\_runlock**, **rms\_wunlock** - kernel reader/writer lock optimized for read-mostly access patterns

**SYNOPSIS**

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

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

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

*void*

```
rm_init(struct rmlock *rm, const char *name);
```

*void*

```
rm_init_flags(struct rmlock *rm, const char *name, int opts);
```

*void*

```
rm_destroy(struct rmlock *rm);
```

*void*

```
rm_rlock(struct rmlock *rm, struct rm_priotracker* tracker);
```

*int*

```
rm_try_rlock(struct rmlock *rm, struct rm_priotracker* tracker);
```

*void*

```
rm_wlock(struct rmlock *rm);
```

*void*

```
rm_runlock(struct rmlock *rm, struct rm_priotracker* tracker);
```

*void*

```
rm_wunlock(struct rmlock *rm);
```

*int*

```
rm_wowned(const struct rmlock *rm);
```

*int*

```
rm_sleep(void *wchan, struct rmlock *rm, int priority, const char *wmesg, int timo);
```

**options INVARIANTS**

**options INVARIANT\_SUPPORT**

*void*

**rm\_assert**(*struct rmlock \*rm, int what*);

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

**RM\_SYSINIT**(*name, struct rmlock \*rm, const char \*desc*);

**RM\_SYSINIT\_FLAGS**(*name, struct rmlock \*rm, const char \*desc, int flags*);

*void*

**rms\_init**(*struct rmslock \*rms, const char \*name*);

*void*

**rms\_destroy**(*struct rmslock \*rms*);

*void*

**rms\_rlock**(*struct rmslock \*rms*);

*void*

**rms\_wlock**(*struct rmslock \*rms*);

*void*

**rms\_runlock**(*struct rmslock \*rms*);

*void*

**rms\_wunlock**(*struct rmslock \*rms*);

## DESCRIPTION

Read-mostly locks allow shared access to protected data by multiple threads, or exclusive access by a single thread. The threads with shared access are known as *readers* since they only read the protected data. A thread with exclusive access is known as a *writer* since it can modify protected data.

Read-mostly locks are designed to be efficient for locks almost exclusively used as reader locks and as such should be used for protecting data that rarely changes. Acquiring an exclusive lock after the lock has been locked for shared access is an expensive operation.

Normal read-mostly locks are similar to `rwlock(9)` locks and follow the same lock ordering rules as `rwlock(9)` locks. Read-mostly locks have full priority propagation like mutexes. Unlike `rwlock(9)`,

read-mostly locks propagate priority to both readers and writers. This is implemented via the *rm\_priotracker* structure argument supplied to **rm\_rlock()** and **rm\_runlock()**. Readers can recurse if the lock is initialized with the RM\_RECURSE option; however, writers are never allowed to recurse.

Sleeping for writers can be allowed by passing RM\_SLEEPABLE to **rm\_init\_flags()**. It changes lock ordering rules to the same as for sx(9) locks. They do not propagate priority to writers, but they do propagate priority to readers. Note that readers are not permitted to sleep regardless of the flag.

Sleepable read-mostly locks (created with **rms\_init()**) allow sleeping for both readers and writers, but don't do priority propagation for either. They follow sx(9) lock ordering.

### Macros and Functions

**rm\_init**(*struct rmlock \*rm, const char \*name*)

Initialize the read-mostly lock *rm*. The *name* description is used solely for debugging purposes. This function must be called before any other operations on the lock.

**rm\_init\_flags**(*struct rmlock \*rm, const char \*name, int opts*)

Similar to **rm\_init()**, initialize the read-mostly lock *rm* with a set of optional flags. The *opts* arguments contains one or more of the following flags:

RM\_NOWITNESS Instruct witness(4) to ignore this lock.

RM\_RECURSE Allow threads to recursively acquire shared locks for *rm*.

RM\_SLEEPABLE Create a sleepable read-mostly lock.

RM\_NEW If the kernel has been compiled with **option INVARIANTS**, **rm\_init\_flags()** will assert that the *rm* has not been initialized multiple times without intervening calls to **rm\_destroy()** unless this option is specified.

RM\_DUPOK witness(4) should not log messages about duplicate locks being acquired.

**rm\_rlock**(*struct rmlock \*rm, struct rm\_priotracker\* tracker*)

Lock *rm* as a reader using *tracker* to track read owners of a lock for priority propagation. This data structure is only used internally by **rmlock** and must persist until **rm\_runlock()** has been called. This data structure can be allocated on the stack since readers cannot sleep. If any thread holds this lock exclusively, the current thread blocks, and its priority is propagated to the exclusive holder. If the lock was initialized with the RM\_RECURSE option the **rm\_rlock()** function can be called when the current thread has already acquired reader access on *rm*.

**rm\_try\_rlock**(*struct rmlock \*rm, struct rm\_priotracker\* tracker*)

Try to lock *rm* as a reader. **rm\_try\_rlock**() will return 0 if the lock cannot be acquired immediately; otherwise, the lock will be acquired and a non-zero value will be returned. Note that **rm\_try\_rlock**() may fail even while the lock is not currently held by a writer. If the lock was initialized with the `RM_RECURSE` option, **rm\_try\_rlock**() will succeed if the current thread has already acquired reader access.

**rm\_wlock**(*struct rmlock \*rm*)

Lock *rm* as a writer. If there are any shared owners of the lock, the current thread blocks. The **rm\_wlock**() function cannot be called recursively.

**rm\_runlock**(*struct rmlock \*rm, struct rm\_priotracker\* tracker*)

This function releases a shared lock previously acquired by **rm\_rlock**(). The *tracker* argument must match the *tracker* argument used for acquiring the shared lock

**rm\_wunlock**(*struct rmlock \*rm*)

This function releases an exclusive lock previously acquired by **rm\_wlock**().

**rm\_destroy**(*struct rmlock \*rm*)

This function destroys a lock previously initialized with **rm\_init**(). The *rm* lock must be unlocked.

**rm\_wowned**(*const struct rmlock \*rm*)

This function returns a non-zero value if the current thread owns an exclusive lock on *rm*.

**rm\_sleep**(*void \*wchan, struct rmlock \*rm, int priority, const char \*wmesg, int timo*)

This function atomically releases *rm* while waiting for an event. The *rm* lock must be exclusively locked. For more details on the parameters to this function, see `sleep(9)`.

**rm\_assert**(*struct rmlock \*rm, int what*)

This function asserts that the *rm* lock is in the state specified by *what*. If the assertions are not true and the kernel is compiled with **options INVARIANTS** and **options INVARIANT\_SUPPORT**, the kernel will panic. Currently the following base assertions are supported:

`RA_LOCKED` Assert that current thread holds either a shared or exclusive lock of *rm*.

`RA_RLOCKED` Assert that current thread holds a shared lock of *rm*.

`RA_WLOCKED` Assert that current thread holds an exclusive lock of *rm*.

**RA\_UNLOCKED** Assert that current thread holds neither a shared nor exclusive lock of *rm*.

In addition, one of the following optional flags may be specified with **RA\_LOCKED**, **RA\_RLOCKED**, or **RA\_WLOCKED**:

**RA\_RECURSED** Assert that the current thread holds a recursive lock of *rm*.

**RA\_NOTRECURSED** Assert that the current thread does not hold a recursive lock of *rm*.

**rms\_init**(*struct rmslock \*rms, const char \*name*)

Initialize the sleepable read-mostly lock *rms*. The *name* description is used as *wmesg* parameter to the `msleep(9)` routine. This function must be called before any other operations on the lock.

**rms\_rlock**(*struct rmlock \*rm*)

Lock *rms* as a reader. If any thread holds this lock exclusively, the current thread blocks.

**rms\_wlock**(*struct rmslock \*rms*)

Lock *rms* as a writer. If the lock is already taken, the current thread blocks. The **rms\_wlock()** function cannot be called recursively.

**rms\_runlock**(*struct rmslock \*rms*)

This function releases a shared lock previously acquired by **rms\_rlock()**.

**rms\_wunlock**(*struct rmslock \*rms*)

This function releases an exclusive lock previously acquired by **rms\_wlock()**.

**rms\_destroy**(*struct rmslock \*rms*)

This functions destroys a lock previously initialized with **rms\_init()**. The *rms* lock must be unlocked.

## SEE ALSO

`locking(9)`, `mutex(9)`, `panic(9)`, `rwlock(9)`, `sema(9)`, `sleep(9)`, `sx(9)`

## HISTORY

These functions appeared in FreeBSD 7.0.

## AUTHORS

The **rmlock** facility was written by Stephan Uphoff. This manual page was written by Gleb Smirnov for `rwlock` and modified to reflect `rmlock` by Stephan Uphoff.

**BUGS**

The **rmlock** implementation is currently not optimized for single processor systems.

**rm\_try\_rlock()** can fail transiently even when there is no writer, while another reader updates the state on the local CPU.

The **rmlock** implementation uses a single per CPU list shared by all rmllocks in the system. If rmllocks become popular, hashing to multiple per CPU queues may be needed to speed up the writer lock process.