

**NAME**

**cap\_new, cap\_getrights** — System calls to manipulate capabilities

**LIBRARY**

Standard C Library (libc, -lc)

**SYNOPSIS**

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

int
cap_new(int fd, cap_rights_t rights);

int
cap_getrights(int fd, cap_rights_t *rightsp);
```

**DESCRIPTION**

Capabilities are special file descriptors derived from an existing file descriptor, such as one returned by `fhopen(2)`, `kqueue(2)`, `mqueue(2)`, `open(2)`, `pipe(2)`, `shm_open(2)`, `socket(2)`, or `socketpair(2)`, but with a restricted set of permitted operations determined by a rights mask set when the capability is created. These restricted rights cannot be changed after the capability is created, although further capabilities with yet more restricted rights may be created from an existing capability. In every other sense, a capability behaves in the same way as the file descriptor it was created from.

`cap_new()` creates a new capability for the existing file descriptor `fd`, and returns a file descriptor for it. Operations on the capability will be limited to those permitted by `rights`, which is static for the lifetime of the capability. If `fd` refers to an existing capability, then `rights` must be equal to or a subset of the rights on that capability. As with `dup(2)` and `dup2(2)`, many properties are shared between the new capability and the existing file descriptor, including open file flags, blocking disposition, and file offset. Many applications will prefer to use the `cap_limitfd(3)` library call, part of `libcapsicum(3)`, as it offers a more convenient interface.

`cap_getrights()` queries the rights associated with the capability referred to by file descriptor `fd`.

These system calls, when combined with `cap_enter(2)`, may be used to construct process sandboxes with highly granular rights assignment.

**RIGHTS**

The following rights may be specified in a new capability rights mask:

CAP_ACCEPT	Permit <code>accept(2)</code> .
CAP_ACL_CHECK	Permit checking of an ACL on a file descriptor; there is no cross-reference for this system call.
CAP_ACL_DELETE	Permit <code>acl_delete_fd_np(3)</code> .
CAP_ACL_GET	Permit <code>acl_get_fd(3)</code> and <code>acl_get_fd_np(3)</code> .
CAP_ACL_SET	Permit <code>acl_set_fd(3)</code> and <code>acl_set_fd_np(3)</code> .
CAP_BIND	Permit <code>bind(2)</code> . Note that sockets can also become bound implicitly as a result of <code>connect(2)</code> or <code>send(2)</code> , and that socket options set with <code>setsockopt(2)</code> may also affect binding behavior.
CAP_CONNECT	Permit <code>connect(2)</code> ; also required for <code>sendto(2)</code> with a non-NULL destination address.

CAP_EVENT	Permit <code>select(2)</code> , <code>poll(2)</code> , and <code>kevent(2)</code> to be used in monitoring the file descriptor for events.
CAP_FEXECVE	Permit <code>fexecve(2)</code> ; <code>CAP_READ</code> will also be required.
CAP_EXTATTR_DELETE	Permit <code>extattr_delete_fd(2)</code> .
CAP_EXTATTR_GET	Permit <code>extattr_get_fd(2)</code> .
CAP_EXTATTR_LIST	Permit <code>extattr_list_fd(2)</code> .
CAP_EXTATTR_SET	Permit <code>extattr_set_fd(2)</code> .
CAP_FCHDIR	Permit <code>fchdir(2)</code> .
CAP_FCHFLAGS	Permit <code>fchflags(2)</code> .
CAP_FCHMOD	Permit <code>fchmod(2)</code> .
CAP_FCHOWN	Permit <code>fchown(2)</code> .
CAP_FCNTL	Permit <code>fcntl(2)</code> ; be aware that this call provides indirect access to other operations, such as <code>flock(2)</code> .
CAP_FLOCK	Permit <code>flock(2)</code> and related calls.
CAP_FPATHCONF	Permit <code>fpathconf(2)</code> .
CAP_FSCK	Permit UFS background-fsck operations on the descriptor.
CAP_FSTAT	Permit <code>fstat(2)</code> .
CAP_FSTATFS	Permit <code>fstatfs(2)</code> .
CAP_FSYNC	Permit <code>aio_fsync(2)</code> and <code>fsync(2)</code> .
CAP_FTRUNCATE	Permit <code>ftruncate(2)</code> .
CAP_FUTIMES	Permit <code>futimes(2)</code> .
CAP_GETPEERNAME	Permit <code>getpeername(2)</code> .
CAP_GETSOCKNAME	Permit <code>getsockname(2)</code> .
CAP_GETSOCKOPT	Permit <code>getsockopt(2)</code> .
CAP_IOCTL	Permit <code>ioctl(2)</code> . Be aware that this system call has enormous scope, including potentially global scope for some objects.
CAP_KEVENT	Permit <code>kevent(2)</code> ; <code>CAP_EVENT</code> is also required on file descriptors that will be monitored using <code>kevent(2)</code> .
CAP_LISTEN	Permit <code>listen(2)</code> ; not much use (generally) without <code>CAP_BIND</code> .
CAP_LOOKUP	Permit the file descriptor to be used as a starting directory for calls such as <code>linkat(2)</code> , <code>openat(2)</code> , and <code>unlinkat(2)</code> . Note that these calls are not available in capability mode as they manipulate a global name space; see <code>cap_enter(2)</code> for details.
CAP_MAC_GET	Permit <code>mac_get_fd(3)</code> .
CAP_MAC_SET	Permit <code>mac_set_fd(3)</code> .
CAP_MMAP	Permit <code>mmap(2)</code> ; specific invocations may also require <code>CAP_READ</code> or <code>CAP_WRITE</code> .

CAP_PDGETPID	Permit <code>pdgetpid(2)</code> .
CAP_PDKILL	Permit <code>pdkill(2)</code> .
CAP_PDWAIT	Permit <code>pdwait4(2)</code> .
CAP_PEELOFF	Permit <code>sctp_peeloff(2)</code> .
CAP_READ	Allow <code>aio_read(2)</code> , <code>pread(2)</code> , <code>read(2)</code> , <code>recv(2)</code> , <code>recvfrom(2)</code> , <code>recvmsg(2)</code> , and related system calls.  For files and other seekable objects, <code>CAP_SEEK</code> may also be required.
CAP_REVOKE	Permit <code>frevoke(2)</code> in certain ABI compatibility modes that support this system call.
CAP_SEEK	Permit operations that seek on the file descriptor, such as <code>lseek(2)</code> , but also required for I/O system calls that modify the file offset, such as <code>read(2)</code> and <code>write(2)</code> .
CAP_SEM_GETVALUE	Permit <code>sem_getvalue(3)</code> .
CAP_SEM_POST	Permit <code>sem_post(3)</code> .
CAP_SEM_WAIT	Permit <code>sem_wait(3)</code> and <code>sem_trywait(3)</code> .
CAP_SETSOCKOPT	Permit <code>setsockopt(2)</code> ; this controls various aspects of socket behavior and may affect binding, connecting, and other behaviors with global scope.
CAP_SHUTDOWN	Permit explicit <code>shutdown(2)</code> ; closing the socket will also generally shut down any connections on it.
CAP_TTYHOOK	Allow configuration of TTY hooks, such as <code>snp(4)</code> , on the file descriptor.
CAP_WRITE	Allow <code>aio_write(2)</code> , <code>pwrite(2)</code> , <code>send(2)</code> , <code>sendmsg(2)</code> , <code>sendto(2)</code> , <code>write(2)</code> , and related system calls.  For files and other seekable objects, <code>CAP_SEEK</code> may also be required.  For <code>sendto(2)</code> with a non-NULL connection address, <code>CAP_CONNECT</code> is also required.

**CAVEAT**

The `cap_new()` system call and the capabilities it creates may be used to assign fine-grained rights to sandboxed processes running in capability mode. However, the semantics of objects accessed via file descriptors are complex, so caution should be exercised in passing object capabilities into sandboxes.

**RETURN VALUES**

If successful, `cap_new()` returns a non-negative integer, termed a file descriptor. It returns -1 on failure, and sets `errno` to indicate the error.

The `cap_getrights()` function returns the value 0 if successful; otherwise the value -1 is returned and the global variable `errno` is set to indicate the error.

**ERRORS**

`cap_new()` may return the following errors:

[EBADF]           The *fd* argument is not a valid active descriptor.

- [EINVAL] An invalid right has been requested in *rights*.
- [EMFILE] The process has already reached its limit for open file descriptors.
- [ENFILE] The system file table is full.
- [EPERM] *rights* contains requested rights not present in the current rights mask associated with the capability referenced by *fd*, if any.

**cap\_getrights()** may return the following errors:

- [EBADF] The *fd* argument is not a valid active descriptor.
- [EINVAL] The *fd* argument is not a capability.

### SEE ALSO

accept(2), aio\_fsync(2), aio\_read(2), aio\_write(2), bind(2), cap\_enter(2), connect(2), dup(2), dup2(2), extattr\_delete\_fd(2), extattr\_get\_fd(2), extattr\_list\_fd(2), extattr\_set\_fd(2), fchflags(2), fchown(2), fcntl(2), fexecve(2), fhopen(2), flock(2), fpathconf(2), fstat(2), fstatfs(2), fsync(2), ftruncate(2), futimes(2), getpeername(2), getsockname(2), getsockopt(2), ioctl(2), kevent(2), kqueue(2), linkat(2), listen(2), mmap(2), mq\_open(2), open(2), openat(2), pdgetpid(2), pdkill(2), pdwait4(2), pipe(2), poll(2), pread(2), pwrite(2), read(2), recv(2), recvfrom(2), recvmsg(2), sctp\_peeloff(2), select(2), send(2), sendmsg(2), sendto(2), setsockopt(2), shm\_open(2), shutdown(2), socket(2), socketpair(2), unlinkat(2), write(2), acl\_delete\_fd\_np(3), acl\_get\_fd(3), acl\_get\_fd\_np(3), acl\_set\_fd\_np(3), cap\_limitfd(3), libcapsicum(3), mac\_get\_fd(3), mac\_set\_fd(3), sem\_getvalue(3), sem\_post(3), sem\_trywait(3), sem\_wait(3), capsicum(4), snp(4)

### HISTORY

Support for capabilities and capabilities mode was developed as part of the TrustedBSD Project.

### AUTHORS

These functions and the capability facility were created by Robert N. M. Watson at the University of Cambridge Computer Laboratory with support from a grant from Google, Inc.

### BUGS

This man page should list the set of permitted system calls more specifically for each capability right.

Capability rights sometimes have unclear indirect impacts, which should be documented, or at least hinted at.