

NAME

Capability rights - Capsicum capability rights for file descriptors

DESCRIPTION

When a file descriptor is created by a function such as `accept(2)`, `accept4(2)`, `fhopen(2)`, `kqueue(2)`, `mq_open(2)`, `open(2)`, `openat(2)`, `pdfork(2)`, `pipe(2)`, `shm_open(2)`, `socket(2)` or `socketpair(2)`, it is assigned all capability rights. Those rights can be reduced (but never expanded) by using the `cap_rights_limit(2)`, `cap_fcntls_limit(2)` and `cap_ioctls_limit(2)` system calls. Once capability rights are reduced, operations on the file descriptor will be limited to those permitted by rights.

The complete list of capability rights is provided below. The `cap_rights_t` type is used to store list of capability rights. The `cap_rights_init(3)` family of functions should be used to manage the structure.

RIGHTS

The following rights may be specified in a rights mask:

<code>CAP_ACCEPT</code>	Permit <code>accept(2)</code> and <code>accept4(2)</code> .
<code>CAP_ACL_CHECK</code>	Permit <code>acl_valid_fd_np(3)</code> .
<code>CAP_ACL_DELETE</code>	Permit <code>acl_delete_fd_np(3)</code> .
<code>CAP_ACL_GET</code>	Permit <code>acl_get_fd(3)</code> and <code>acl_get_fd_np(3)</code> .
<code>CAP_ACL_SET</code>	Permit <code>acl_set_fd(3)</code> and <code>acl_set_fd_np(3)</code> .
<code>CAP_BIND</code>	When not in capabilities mode, permit <code>bind(2)</code> and <code>bindat(2)</code> with special value <code>AT_FDCWD</code> in the <code>fd</code> parameter. 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.
<code>CAP_BINDAT</code>	Permit <code>bindat(2)</code> . This right has to be present on the directory descriptor. This right includes the <code>CAP_LOOKUP</code> right.
<code>CAP_CHFLAGSAT</code>	An alias to <code>CAP_FCHFLAGS</code> and <code>CAP_LOOKUP</code> .
<code>CAP_CONNECT</code>	When not in capabilities mode, permit <code>connect(2)</code> and <code>connectat(2)</code> with special value <code>AT_FDCWD</code> in the <code>fd</code> parameter. This right is also required for <code>sendto(2)</code> with a non-NULL destination address.

CAP_CONNECTAT	Permit connectat(2). This right has to be present on the directory descriptor. This right includes the CAP_LOOKUP right.
CAP_CREATE	Permit openat(2) with the O_CREAT flag.
CAP_EVENT	Permit select(2), poll(2), and kevent(2) to be used in monitoring the file descriptor for events.
CAP_EXTATTR_DELETE	Permit extattr_delete_fd(2).
CAP_EXTATTR_GET	Permit extattr_get_fd(2).
CAP_EXTATTR_LIST	Permit extattr_list_fd(2).
CAP_EXTATTR_SET	Permit extattr_set_fd(2).
CAP_FCHDIR	Permit fchdir(2).
CAP_FCHFLAGS	Permit fchflags(2) and chflagsat(2) if the CAP_LOOKUP right is also present.
CAP_FCHMOD	Permit fchmod(2) and fchmodat(2) if the CAP_LOOKUP right is also present.
CAP_FCHMODAT	An alias to CAP_FCHMOD and CAP_LOOKUP.
CAP_FCHOWN	Permit fchown(2) and fchownat(2) if the CAP_LOOKUP right is also present.
CAP_FCHOWNAT	An alias to CAP_FCHOWN and CAP_LOOKUP.
CAP_FCNTL	Permit fcntl(2). Note that only the F_GETFL, F_SETFL, F_GETOWN and F_SETOWN commands require this capability right. Also note that the list of permitted commands can be further limited with the cap_fcntls_limit(2) system call.
CAP_FEXECVE	Permit fexecve(2) and openat(2) with the O_EXEC flag; CAP_READ is also required.
CAP_FLOCK	Permit flock(2), fcntl(2) (with F_GETLK, F_SETLK, F_SETLKW or

	F_SETLK_REMOTE flag) and <code>openat(2)</code> (with <code>O_EXLOCK</code> or <code>O_SHLOCK</code> flag).
<code>CAP_FPATHCONF</code>	Permit <code>fpathconf(2)</code> .
<code>CAP_FSCK</code>	Permit UFS background-fsck operations on the descriptor.
<code>CAP_FSTAT</code>	Permit <code>fstat(2)</code> and <code>fstatat(2)</code> if the <code>CAP_LOOKUP</code> right is also present.
<code>CAP_FSTATAT</code>	An alias to <code>CAP_FSTAT</code> and <code>CAP_LOOKUP</code> .
<code>CAP_FSTATFS</code>	Permit <code>fstatfs(2)</code> .
<code>CAP_FSYNC</code>	Permit <code>aio_fsync(2)</code> , <code>fdatasync(2)</code> , <code>fsync(2)</code> and <code>openat(2)</code> with <code>O_FSYNC</code> or <code>O_SYNC</code> flag.
<code>CAP_FTRUNCATE</code>	Permit <code>ftruncate(2)</code> and <code>openat(2)</code> with the <code>O_TRUNC</code> flag.
<code>CAP_FUTIMES</code>	Permit <code>futimens(2)</code> and <code>futimes(2)</code> , and permit <code>futimesat(2)</code> and <code>utimensat(2)</code> if the <code>CAP_LOOKUP</code> right is also present.
<code>CAP_FUTIMESAT</code>	An alias to <code>CAP_FUTIMES</code> and <code>CAP_LOOKUP</code> .
<code>CAP_GETPEERNAME</code>	Permit <code>getpeername(2)</code> .
<code>CAP_GETSOCKNAME</code>	Permit <code>getsockname(2)</code> .
<code>CAP_GETSOCKOPT</code>	Permit <code>getsockopt(2)</code> .
<code>CAP_IOCTL</code>	Permit <code>ioctl(2)</code> . Be aware that this system call has enormous scope, including potentially global scope for some objects. The list of permitted <code>ioctl</code> commands can be further limited with the <code>cap_ioctls_limit(2)</code> system call.
<code>CAP_KQUEUE</code>	An alias to <code>CAP_KQUEUE_CHANGE</code> and <code>CAP_KQUEUE_EVENT</code> .
<code>CAP_KQUEUE_CHANGE</code>	Permit <code>kevent(2)</code> on a <code>kqueue(2)</code> descriptor that modifies list of monitored events (the <i>changelist</i> argument is non-NULL).
<code>CAP_KQUEUE_EVENT</code>	Permit <code>kevent(2)</code> on a <code>kqueue(2)</code> descriptor that monitors events (the

eventlist argument is non-NULL). `CAP_EVENT` is also required on file descriptors that will be monitored using `kevent(2)`.

<code>CAP_LINKAT_SOURCE</code>	Permit <code>linkat(2)</code> on the source directory descriptor. This right includes the <code>CAP_LOOKUP</code> right. Warning: <code>CAP_LINKAT_SOURCE</code> makes it possible to link files in a directory for which file descriptors exist that have additional rights. For example, a file stored in a directory that does not allow <code>CAP_READ</code> may be linked in another directory that does allow <code>CAP_READ</code> , thereby granting read access to a file that is otherwise unreadable.
<code>CAP_LINKAT_TARGET</code>	Permit <code>linkat(2)</code> on the target directory descriptor. This right includes the <code>CAP_LOOKUP</code> right.
<code>CAP_LISTEN</code>	Permit <code>listen(2)</code> ; not much use (generally) without <code>CAP_BIND</code> .
<code>CAP_LOOKUP</code>	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> .
<code>CAP_MAC_GET</code>	Permit <code>mac_get_fd(3)</code> .
<code>CAP_MAC_SET</code>	Permit <code>mac_set_fd(3)</code> .
<code>CAP_MKDIRAT</code>	Permit <code>mkdirat(2)</code> . This right includes the <code>CAP_LOOKUP</code> right.
<code>CAP_MKFIFOAT</code>	Permit <code>mkfifoat(2)</code> . This right includes the <code>CAP_LOOKUP</code> right.
<code>CAP_MKNODAT</code>	Permit <code>mknodat(2)</code> . This right includes the <code>CAP_LOOKUP</code> right.
<code>CAP_MMAP</code>	Permit <code>mmap(2)</code> with the <code>PROT_NONE</code> protection.
<code>CAP_MMAP_R</code>	Permit <code>mmap(2)</code> with the <code>PROT_READ</code> protection. This right includes the <code>CAP_READ</code> and <code>CAP_SEEK</code> rights.
<code>CAP_MMAP_RW</code>	An alias to <code>CAP_MMAP_R</code> and <code>CAP_MMAP_W</code> .
<code>CAP_MMAP_RWX</code>	An alias to <code>CAP_MMAP_R</code> , <code>CAP_MMAP_W</code> and <code>CAP_MMAP_X</code> .
<code>CAP_MMAP_RX</code>	An alias to <code>CAP_MMAP_R</code> and <code>CAP_MMAP_X</code> .

CAP_MMAP_W	Permit <code>mmap(2)</code> with the <code>PROT_WRITE</code> protection. This right includes the <code>CAP_WRITE</code> and <code>CAP_SEEK</code> rights.
CAP_MMAP_WX	An alias to <code>CAP_MMAP_W</code> and <code>CAP_MMAP_X</code> .
CAP_MMAP_X	Permit <code>mmap(2)</code> with the <code>PROT_EXEC</code> protection. This right includes the <code>CAP_SEEK</code> right.
CAP_PDGETPID	Permit <code>pdgetpid(2)</code> .
CAP_PDKILL	Permit <code>pdkill(2)</code> .
CAP_PEELOFF	Permit <code>sctp_peeloff(2)</code> .
CAP_PREAD	An alias to <code>CAP_READ</code> and <code>CAP_SEEK</code> .
CAP_PWRITE	An alias to <code>CAP_SEEK</code> and <code>CAP_WRITE</code> .
CAP_READ	Permit <code>aio_read(2)</code> (<code>CAP_SEEK</code> is also required), <code>openat(2)</code> with the <code>O_RDONLY</code> flag, <code>read(2)</code> , <code>readv(2)</code> , <code>recv(2)</code> , <code>recvfrom(2)</code> , <code>recvmsg(2)</code> , <code>pread(2)</code> (<code>CAP_SEEK</code> is also required), <code>preadv(2)</code> (<code>CAP_SEEK</code> is also required) and related system calls.
CAP_RECV	An alias to <code>CAP_READ</code> .
CAP_RENAMEAT_SOURCE	Permit <code>renameat(2)</code> on the source directory descriptor. This right includes the <code>CAP_LOOKUP</code> right. Warning: <code>CAP_RENAMEAT_SOURCE</code> makes it possible to move files to a directory for which file descriptors exist that have additional rights. For example, a file stored in a directory that does not allow <code>CAP_READ</code> may be moved to another directory that does allow <code>CAP_READ</code> , thereby granting read access to a file that is otherwise unreadable.
CAP_RENAMEAT_TARGET	Permit <code>renameat(2)</code> on the target directory descriptor. This right includes the <code>CAP_LOOKUP</code> right.
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 can read or write at any position in the file, such as <code>pread(2)</code> and <code>pwrite(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_SEND	An alias to <code>CAP_WRITE</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 <code>explicit_shutdown(2)</code> ; closing the socket will also generally shut down any connections on it.
CAP_SYMLINKAT	Permit <code>symlinkat(2)</code> . This right includes the <code>CAP_LOOKUP</code> right.
CAP_TTYHOOK	Allow configuration of TTY hooks, such as <code>snp(4)</code> , on the file descriptor.
CAP_UNLINKAT	Permit <code>unlinkat(2)</code> and <code>renameat(2)</code> . This right is only required for <code>renameat(2)</code> on the destination directory descriptor if the destination object already exists and will be removed by the <code>rename</code> . This right includes the <code>CAP_LOOKUP</code> right.
CAP_WRITE	Allow <code>aio_write(2)</code> , <code>openat(2)</code> with <code>O_WRONLY</code> and <code>O_APPEND</code> flags set, <code>send(2)</code> , <code>sendmsg(2)</code> , <code>sendto(2)</code> , <code>write(2)</code> , <code>writev(2)</code> , <code>pwrite(2)</code> , <code>pwritev(2)</code> and related system calls. For <code>sendto(2)</code> with a non-NULL connection address, <code>CAP_CONNECT</code> is also required. For <code>openat(2)</code> with the <code>O_WRONLY</code> flag, but without the <code>O_APPEND</code> flag, <code>CAP_SEEK</code> is also required. For <code>aio_write(2)</code> , <code>pwrite(2)</code> and <code>pwritev(2)</code> <code>CAP_SEEK</code> is also required.

SEE ALSO

`accept(2)`, `accept4(2)`, `aio_fsync(2)`, `aio_read(2)`, `aio_write(2)`, `bind(2)`, `bindat(2)`, `cap_enter(2)`, `cap_fcntls_limit(2)`, `cap_ioctls_limit(2)`, `cap_rights_limit(2)`, `chflagsat(2)`, `connect(2)`, `connectat(2)`, `extattr_delete_fd(2)`, `extattr_get_fd(2)`, `extattr_list_fd(2)`, `extattr_set_fd(2)`, `fchflags(2)`, `fchmod(2)`, `fchmodat(2)`, `fchown(2)`, `fchownat(2)`, `fcntl(2)`, `fexecve(2)`, `fhopen(2)`, `flock(2)`, `fpathconf(2)`, `fstat(2)`, `fstatat(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)`, `pdfork(2)`,

pdgetpid(2), pdkill(2), pdwait4(2), pipe(2), poll(2), pread(2), preadv(2), pwrite(2), pwritev(2), read(2), readv(2), recv(2), recvfrom(2), recvmsg(2), renameat(2), sctp_peeloff(2), select(2), send(2), sendmsg(2), sendto(2), setsockopt(2), shm_open(2), shutdown(2), socket(2), socketpair(2), symlinkat(2), unlinkat(2), write(2), writev(2), acl_delete_fd_np(3), acl_get_fd(3), acl_get_fd_np(3), acl_set_fd(3), acl_set_fd_np(3), acl_valid_fd_np(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

This manual page was created by Pawel Jakub Dawidek <pawel@dawidek.net> under sponsorship from the FreeBSD Foundation based on the `cap_new(2)` manual page by Robert Watson <rwatson@FreeBSD.org>.