

**NAME**

capsh - capability shell wrapper

**SYNOPSIS**

**capsh** [*OPTION*]...

**DESCRIPTION**

Linux capability support and use can be explored and constrained with this tool. This tool provides a handy wrapper for certain types of capability testing and environment creation. It also provides some debugging features useful for summarizing capability state.

**OPTIONS**

The tool takes a number of optional arguments, acting on them in the order they are provided. They are as follows:

<b>--print</b>	Display prevailing capability and related state.
<b>-- [args]</b>	Execute <b>/bin/bash</b> with trailing arguments. Note, you can use <b>-c 'command to execute'</b> for specific commands.
<b>==</b>	Execute <b>capsh</b> again with remaining arguments. Useful for testing <b>exec()</b> behavior.
<b>--caps=cap-set</b>	Set the prevailing process capabilities to those specified by <i>cap-set</i> . Where <i>cap-set</i> is a text-representation of capability state as per <b>cap_from_text(3)</b> .
<b>--drop=cap-list</b>	Remove the listed capabilities from the prevailing bounding set. The capabilities are a comma separated list of capabilities as recognized by the <b>cap_from_name(3)</b> function. Use of this feature requires that the capsh program is operating with <b>CAP_SETPCAP</b> in its effective set.
<b>--inh=cap-list</b>	Set the inheritable set of capabilities for the current process to equal those provided in the comma separated list. For this action to succeed, the prevailing process should already have each of these capabilities in the union of the current inheritable and permitted capability sets, or the capsh program is operating with <b>CAP_SETPCAP</b> in its effective set.
<b>--user=username</b>	Assume the identity of the named user. That is, look up the user's <i>uid</i> and <i>gid</i> with <b>getpwuid(3)</b> and their group memberships with <b>getgrouplist(3)</b> and set them all.
<b>--uid=id</b>	Force all <b>uid</b> values to equal <i>id</i> using the <b>setuid(2)</b> system call.
<b>--gid=&lt;id&gt;</b>	Force all <b>gid</b> values to equal <i>id</i> using the <b>setgid(2)</b> system call.
<b>--groups=&lt;id-list&gt;</b>	Set the supplementary groups to the numerical list provided. The groups are set with the <b>setgroups(2)</b> system call.
<b>--keep=&lt;0/1&gt;</b>	In a non-pure capability mode, the kernel provides liberal privilege to the super-user. However, it is normally the case that when the super-user changes <i>uid</i> to some lesser user, then capabilities are dropped. For these situations, the kernel can permit the process to retain its capabilities after a <b>setuid(2)</b> system call. This feature is known as <i>keep-caps</i> support. The way to activate it using this script is with this argument. Setting the value to 1 will cause <i>keep-caps</i> to be active. Setting it to 0 will cause <i>keep-caps</i> to deactivate for the current process. In all cases, <i>keep-caps</i> is deactivated when an <b>exec()</b> is performed. See <b>--secbits</b> for ways to disable this feature.
<b>--secbits=N</b>	XXX - need to document this feature.
<b>--chroot=path</b>	Execute the <b>chroot(2)</b> system call with the new root-directory ( <i>/</i> ) equal to <i>path</i> . This operation requires <b>CAP_SYS_CHROOT</b> to be in effect.
<b>--forkfor=sec</b>	

**--killit=***sig*

**--decode=***N*

This is a convenience feature. If you look at `/proc/1/status` there are some capability related fields of the following form:

```
CapInh: 0000000000000000  CapPrm: ffffffff CapEff: ffffffffef
CapBnd: ffffffff
```

This option provides a quick way to decode a capability vector represented in this form. For example, the missing capability from this effective set is 0x0100. By running:

```
capsh --decode=0x0100
```

we observe that the missing capability is: **cap\_setpcap**.

**--supports=***xxx*

As the kernel evolves, more capabilities are added. This option can be used to verify the existence of a capability on the system. For example, **--supports=cap\_syslog** will cause capsh to promptly exit with a status of 1 when run on kernel 2.6.27. However, when run on kernel 2.6.38 it will silently succeed.

## EXIT STATUS

Following successful execution the tool exits with status 0. Following an error, the tool immediately exits with status 1.

## AUTHOR

Written by Andrew G. Morgan <morgan@kernel.org>.

## REPORTING BUGS

Please report bugs to the author.

## SEE ALSO

[libcap\(3\)](#), [getcap\(8\)](#), [setcap\(8\)](#) and [capabilities\(7\)](#).