
Subject: Capabilities issue

Posted by [kevinm](#) on Tue, 31 Aug 2010 14:35:44 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi All !

I appear to be having an issue with capabilities inside a openvz container..

The source code that is giving me issues is the following :

```
/* init cap with all zeros */
cap = cap_init();
capval[0] = CAP_SETUID;
capval[1] = CAP_SETGID;
capval[2] = CAP_DAC_READ_SEARCH;
capval[3] = CAP_SYS_CHROOT;
cap_set_flag(cap, CAP_PERMITTED, (chroot_root >= 0 ? 4 : 3), capval, CAP_SET);
if (cap_set_proc(cap) != 0) {
    ap_log_error(APLOG_MARK, APLOG_ERR, 0, NULL, "%s CRITICAL ERROR
ruid_child_init:cap_set_proc failed", MODULE_NAME);
}
cap_free(cap);
```

I have granted the following capabilities to the virtual server , and restarted it :

Quote:CAPABILITY="CHOWN:on DAC_READ_SEARCH:on DAC_OVERRIDE:on SETGID:on
SETUID:on NET_BIND_SERVICE:on NET_ADMIN:on SYS_CHROOT:on SYS_NICE:on
SYS_CHROOT:on "

however I get logged to apache error logs :

Quote:[Tue Aug 31 09:31:50 2010] [error] mod_ruid CRITICAL ERROR ruid_setup:cap_set_proc
failed

stracing a process, shows

Quote:capset(0x19980330, 0, {CAP_DAC_OVERRIDE|CAP_DAC_READ_SEARCH,
CAP_DAC_OVERRIDE|CAP_SETGID|CAP_SETUID, 0}) = -1 EPERM (Operation not permitted)

is there any reason that even though ive granted these capabilities, that im still receiving -1
EPERM (Operation not permitted) , I cant see anything wrong with the capabilities granted to the
ones that are failing, any advise / assistance would be greatly appreciated.

best regards
Kev

Subject: Re: Capabilities issue
Posted by [maratrus](#) on Thu, 02 Sep 2010 13:10:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

try granting SETPCAP capability as well. Hope this helps.

Subject: Re: Capabilities issue
Posted by [kevinm](#) on Thu, 02 Sep 2010 14:13:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi there Maratrus,

Thanks for the answer, I have tried this, and get

Quote:

```
# vzctl restart 130
Restarting container
Stopping container ...
Container was stopped
Container is unmounted
Starting container ...
Container is mounted
Unable to set capability: Operation not permitted
Unable to set capability
Container start failed
```

I see this in the bugzilla :

http://bugzilla.openvz.org/show_bug.cgi?id=554

Is this possible to set ?

Best Regards

Subject: Re: Capabilities issue

Posted by [maratrus](#) on Fri, 03 Sep 2010 14:50:54 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

thanks for providing this bugreport's link.

I'm at a loss at the moment because there is a contradiction between vzctl utility and a kernel from my point of view.

Look, at the kernel side the following checkings are made

```
asmlinkage long sys_capset(cap_user_header_t header, const cap_user_data_t data)
{
<snip>
if (pid && pid != virt_pid(current) && !capable(CAP_SETPCAP))
    return -EPERM;
<snip>
}
```

As I understand, this piece of code implies that it is possible to use "capset" system call from inside the VE. The only thing that must be made is providing CAP_SETPCAP capability to it. The standard way to do it is via vzctl.

On the other hand, in the bugreport provided by you we can read

Looks like all linux kernels has init started with CAP_SETPCAP explicitly disabled due to security implications, so that's why nobody (including vzctl) can set it.

OK, we can't set it on boot. But if that means that it's impossible to set this capability at all (after init is started) it must be mirrored in vzctl's man page.

So, I would recommend you to ask this question directly in the existing bugreport to get developer's opinion about this situation and the ways to work around it.
