
Subject: ***SOLVED*** Problem starting VPS
Posted by [Phoebus](#) on Wed, 15 Feb 2006 14:49:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

I installed VPS on debian, created template and when I try to run it I get this:

```
# vzctl start 1002
Starting VPS ...
VPS is mounted
Unable to set capability: Operation not permitted
Unable to set capability
VPS start failed
VPS is unmounted
#
```

I tried verbose:

```
Unable to open /usr/lib/vzctl/modules/: No such file or directory
Starting VPS ...
Running: /usr/sbin/vzquota show 1002
Running: /usr/sbin/vzquota on 1002 -r 0 -b 1048676 -B 1153534 -i 200100 -l 220100 -e 0 -n 0 -s 0
Mounting root: /vz/root/1002 /vz/private/1002
VPS is mounted
Unable to set capability: Operation not permitted
Unable to set capability
VPS start failed
Running: /usr/sbin/vzquota off 1002
Running: /usr/sbin/vzquota stat 1002 -f
VPS is unmounted
```

I tried tuning vzquota and it works OK. Any ideas how to fix it?

Subject: Re: Problem starting VPS
Posted by [dim](#) on Wed, 15 Feb 2006 15:22:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

Could you provide `uname -a` and `strace -ff vzctl start 1002` output?

Subject: Re: Problem starting VPS
Posted by [Phoebus](#) on Wed, 15 Feb 2006 15:48:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

uname -a:

```
Linux Nero 2.6.15-025stab014 #3 SMP Wed Feb 15 15:51:37 CET 2006 i686 GNU/Linux
```

strace -ff vzctl start 1002

<http://www.apollonus.net/uploadi/trace.out>

Subject: Re: Problem starting VPS
Posted by [dev](#) on Thu, 16 Feb 2006 11:24:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sorry, what you mean by "I tried tuning vzquota and it works OK"? Have you succeeded to start VPS?

Also, please, attach log from:
strace -f vzctl start 1002

since you attached only one file from your prev strace.

P.S. you can use stable 2.6.8 until the issue is resolved.

Subject: Re: Problem starting VPS
Posted by [dwebwarm](#) on Tue, 28 Feb 2006 09:08:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

I have the same problem starting a vps. Here the strace dumps:

```
strace -ff vzctl start 101
http://dev.hive-net.net/dump.txt
http://dev.hive-net.net/dump.txt.20510
http://dev.hive-net.net/dump.txt.20511
http://dev.hive-net.net/dump.txt.20512
http://dev.hive-net.net/dump.txt.20513
http://dev.hive-net.net/dump.txt.20514
http://dev.hive-net.net/dump.txt.20515
```

```
strace -f vzctl start 101
http://dev.hive-net.net/dump2.txt
```

```
uname -a
Linux server2 2.6.15-025stab014-DWebWarm #5 Tue Feb 28 02:04:30 EET 2006 i686 AMD
Athlon(tm) XP 2400+ GNU/Linux
```

```
cat /proc/version
Linux version 2.6.15-025stab014-DWebWarm (root@server2) (gcc version 3.4.5 (Gentoo 3.4.5,
```

ssp-3.4.5-1.0, pie-8.7.9)) #5 Tue Feb 28 02:04:30 EET 2006

I have problems compiling vzctl 3.0.0.2 too, but the errors I'll post in additional thread later.

Subject: Re: Problem starting VPS
Posted by [dev](#) on Tue, 28 Feb 2006 09:20:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you post also output of /proc/self/status, when you are a root on your node.

Do you experience this problem with any of your VPSs?
Please, you 2.6.8 until it is resolved.

Subject: Re: Problem starting VPS
Posted by [dwebwarm](#) on Tue, 28 Feb 2006 11:37:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

Here the output of cat /proc/self/status

```
server2 ~ # cat /proc/self/status
Name: cat
State: R (running)
SleepAVG: 88%
Tgid: 10461
Pid: 10461
PPid: 2986
TracerPid: 0
Uid: 0 0 0 0
Gid: 0 0 0 0
FDSize: 256
Groups: 0 1 2 3 4 6 10 11 20 26 27
envID: 0
VPid: 10461
VmPeak: 1456 kB
VmSize: 1456 kB
VmLck: 0 kB
VmHWM: 360 kB
VmRSS: 360 kB
VmData: 144 kB
VmStk: 88 kB
VmExe: 16 kB
VmLib: 1180 kB
VmPTE: 12 kB
Threads: 1
```

SigQ: 0/12031
SigPnd: 0000000000000000
ShdPnd: 0000000000000000
SigBlk: 0000000000000000
SigIgn: 0000000000000000
SigCgt: 0000000000000000
CapInh: 0000000000000000
CapPrm: 00000000ffffff
CapEff: 00000000ffffeff
TaskUB: 0
MMUB: 0

the odd thing here is, that I'm able to start VPS on another node using P4 2.4 GHz CPU and SATA devices. I was able to start 5 different OS'es in a VPS too (Gentoo, CentoOS, Debian 3.1, Ubuntu 5.10 and Fedora Core 4). After using the same configuration here nothing happens . It's a bit weird for me now. Maybe the problem is in the hardware and not software related. I mean, that OpenVZ is not working for me on Athlon/Sempron CPU arch, only on Intel based systems.

Subject: Re: Problem starting VPS
Posted by [dev](#) on Tue, 28 Feb 2006 13:40:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

Also, one important question.
is it kernel build by us or your own? looks like your own, yeah?
then I need .config file for it.

Subject: Re: Problem starting VPS
Posted by [dwebwarm](#) on Tue, 28 Feb 2006 13:50:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

All my installations are Gentoo based and I'm using the openvz-sources in the portage tree. As I told you in the last reply one time it's working with Intel processor, and sometimes on other arch not.

Here the link for the .config file:

<http://dev.hive-net.net/.config>

Subject: Re: Problem starting VPS
Posted by [dev](#) on Tue, 28 Feb 2006 14:12:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you try this debug patch please?

```

--- ./kernel/capability.c.capdbg 2006-02-28 16:09:25.000000000 +0300
+++ ./kernel/capability.c 2006-02-28 17:04:29.000000000 +0300
@@ -192,8 +192,12 @@ asmlinkage long sys_capset(cap_user_head
    if (get_user(pid, &header->pid))
        return -EFAULT;

-   if (pid && pid != virt_pid(current) && !capable(CAP_SETPCAP))
+   if (pid && pid != virt_pid(current) && !capable(CAP_SETPCAP)) {
+   printk("capset: (%s): pid = %d, vpid = %d, caps=%08lx\n",
+   current->comm, pid, virt_pid(current),
+   current->cap_effective);
+   return -EPERM;
+ }

    if (copy_from_user(&effective, &data->effective, sizeof(effective)) ||
        copy_from_user(&inheritable, &data->inheritable, sizeof(inheritable))) ||
@@ -235,5 +239,9 @@ out:
    read_unlock(&tasklist_lock);
    spin_unlock(&task_capability_lock);

+   if (ret == -EPERM) {
+   printk("sys_capset: %s:%d, pid=%d\n",
+   current->comm, current->pid, pid);
+   }
    return ret;
}

--- ./security/commoncap.c.capdbg 2006-02-28 16:09:25.000000000 +0300
+++ ./security/commoncap.c 2006-02-28 16:59:39.000000000 +0300
@@ -78,27 +78,49 @@ int cap_capget (struct task_struct *targ
int cap_capset_check (struct task_struct *target, kernel_cap_t *effective,
    kernel_cap_t *inheritable, kernel_cap_t *permitted)
{
+ int ret;
+
+ /* Derived from kernel/capability.c:sys_capset. */
+ /* verify restrictions on target's new Inheritable set */
if (!cap_issubset (*inheritable,
    cap_combine (target->cap_inheritable,
    current->cap_permitted))) {
- return -EPERM;
+ ret = 1;
+ goto oops;
}

/* verify restrictions on target's new Permitted set */
if (!cap_issubset (*permitted,

```

```

    cap_combine (target->cap_permitted,
    current->cap_permitted))) {
- return -EPERM;
+ ret = 2;
+ goto oops;
}

/* verify the _new_Effective_ is a subset of the _new_Permitted_ */
if (!cap_issubset (*effective, *permitted)) {
- return -EPERM;
+ ret = 3;
+ goto oops;
}

return 0;
+
+oops:
+ printk("cap_capset_check: %d (%s:%d), curcaps=%08lx:%08lx:%08lx, target (%s:%d)
caps=%08lx:%08lx:%08lx, set %08lx:%08lx:%08lx\n",
+ ret,
+ current->comm, current->pid,
+ current->cap_effective,
+ current->cap_permitted,
+ current->cap_inheritable,
+ target->comm, target->pid,
+ target->cap_effective,
+ target->cap_permitted,
+ target->cap_inheritable,
+ *effective,
+ *permitted,
+ *inheritable
+ );
+ return -EPERM;
}

```

```

void cap_capset_set (struct task_struct *target, kernel_cap_t *effective,
--- ./security/dummy.c.capdbg 2006-01-03 06:21:10.000000000 +0300
+++ ./security/dummy.c 2006-02-28 17:00:37.000000000 +0300
@@ -56,6 +56,7 @@ static int dummy_capset_check (struct ta
    kernel_cap_t * inheritable,
    kernel_cap_t * permitted)
{
+ printk("dummy_capset_check: %s:%d\n", current->comm, current->pid);
return -EPERM;
}

```

```

--- ./security/selinux/hooks.c.capdbg 2006-02-28 16:09:25.000000000 +0300
+++ ./security/selinux/hooks.c 2006-02-28 17:02:41.000000000 +0300

```

```

@@ -1314,9 +1314,15 @@ static int selinux_capset_check(struct t

    error = secondary_ops->capset_check(target, effective, inheritable, permitted);
    if (error)
-   return error;
+   goto err;
+
+ error = task_has_perm(current, target, PROCESS__SETCAP);
+ if (error)
+   goto err;

- return task_has_perm(current, target, PROCESS__SETCAP);
+ return 0;
+err:
+ printk("selinux_capset_check: %s:%d err=%d\n", current->comm, current->pid, error);
}

static void selinux_capset_set(struct task_struct *target, kernel_cap_t *effective,

```

Also I noted, that you have following config options, different from our config and potentially conflicting with virtualization:

CONFIG_SECURITY=y

CONFIG_AUDIT=y

please turn it off.

Maybe you can check the whole kernel with OVZ config?

Subject: Re: Problem starting VPS

Posted by [dwebwarm](#) on Tue, 28 Feb 2006 14:48:18 GMT

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

Hi,

10x for the reply. I'll already patched the kernel and I'm compiling it at the moment. I'll give you response soon.

Subject: Re: Problem starting VPS

Posted by [dwebwarm](#) on Tue, 28 Feb 2006 15:32:39 GMT

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

After applying the debugging patch and removing the CONFIG_SECURITY=y option, everything seems to function perfectly. I think, that the problem was caused by the kernel option.

Subject: Re: Problem starting VPS
Posted by [robataka](#) on Wed, 09 Aug 2006 17:08:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

fwiw, I did not apply the patch, but did remove the two kernel options and it worked.

Thanks for the tip!
