
Subject: "vzctl stop" - hangs

Posted by [dagr](#) on Wed, 14 Nov 2007 14:20:44 GMT

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

Usually i use kernel 2.6.9 , but now i have to move to 2.6.18 (2.6.18-8.1.15.el5.028stab047.1PAE) and i found out that "vzctl stop" - hangs for exactly 2 mins (start is ok and runs ok)

I found out that it happens only with VEs which use loop mounts (mount -o loop,ro) - i use them for sharing data among VEs.

I made strace during stop (it helped me once when stop hanged cause couldnt find "halt" binary), but wasnt able to find out the reason . Anyway i attached strace.txt .

Also i noticed that it doesnt umount clean loop devices .

On prvious kernel i had in .umount

```
umount /vz/root/$VEID/vps_static
```

```
umount /vz/root/$VEID/vps_libs
```

On new kernel i get error

```
[dagr@ssw2-ca shared_vps]$ date;sudo vzctl stop 1051206;date
```

```
Wed Nov 14 14:05:40 GMT 2007
```

```
Stopping VE ...
```

```
VE was stopped
```

```
umount
```

```
umount: /u04/vz/shared_vps/php5_base/vps_libs_loop: not mounted
```

```
umount: /u04/vz/shared_vps/php5_base/vps_libs_loop: not mounted
```

```
Error executing umount script /etc/vz/conf/1051206.umount
```

```
Wed Nov 14 14:07:41 GMT 2007
```

And i get 1 loop device not detached (from /vz/root/\$VEID/vps_libs mount)

```
[dagr@ssw2-ca private]$ ps -ef| grep loop
```

```
root    601    1  0 14:04 ?        00:00:00 [loop1]
```

```
[dagr@ssw2-ca private]$ sudo losetup -d /dev/loop1
```

I thought that may be openvz on new kernel somehow handles loop umounts by itself , but if comment .umount file , it leaves all 2 mounts.

(vzctl-3.0.18-1 , ovzkernel-PAE-2.6.18-8.1.15.el5.028stab047.1)

So 2 questions.

1. Whats wrong with stop ?

2. Whats wrong with loop device umount ?

File Attachments

1) [strace.txt](#), downloaded 305 times

Subject: Re: "vzctl stop" - hangs

Posted by [vaverin](#) on Thu, 15 Nov 2007 07:36:24 GMT

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

1) 2 minutes is "vzctl stop" timeout, it waits until VE will be stopped completely. Obviously initscripts inside your VE are not finished correctly.

Some time ago we have disabled ability to kill init process from inside VE. Now VE's initscripts should call halt or reboot (or any another application that calls sys_reboot syscall) to kill VE's init process and stop it completely.

2) umount problem

Are you sure that external mounts are leaves really? It should be umounted automaticly on VE stop.

Please take a look at the following Vasily Tarasov's comment:

http://forum.openvz.org/index.php?t=msg&goto=13453&&srch=mount+%2Fetc%2Fmtab#msg_13453:

"The thing is that when you produce mount command on HN (e.g. from .mount script), the entry in /etc/mtab on HN appears. Later when VE stops all mount points in VE are umounted automatically even _without_ umount script. But the entry in /etc/mtab is still there! It makes mount a bit crazy...

So, you should use `mount -n` in mount script and don't use umount script at all."

Subject: Re: "vzctl stop" - hangs

Posted by [vaverin](#) on Thu, 15 Nov 2007 07:40:38 GMT

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

Btw, why do not you use bindmounts instead of loop mounts?

Subject: Re: "vzctl stop" - hangs

Posted by [dagr](#) on Thu, 15 Nov 2007 08:22:03 GMT

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

cause i need them to be read only

BTW - mount -n helps with error from mtab , but loop devices are still not detached

```
[dagr@ssw2-ca ~]$ date;sudo vzctl stop 1051206;date
Thu Nov 15 08:29:05 GMT 2007
Stopping VE ...
VE was stopped
umount
VE is unmounted
Thu Nov 15 08:31:06 GMT 2007
[dagr@ssw2-ca ~]$ ps -ef| grep loop
root    25193    1 0 08:28 ?        00:00:00 [loop0]
root    25200    1 0 08:28 ?        00:00:00 [loop1]
dagr    27545 24812  0 08:31 pts/3    00:00:00 grep loop
[dagr@ssw2-ca ~]$ sudo vzctl start 1051206
Starting VE ...
VE is mounted
Adding IP address(es): 10.5.1.206
Setting CPU units: 1000
Configure meminfo: 149152
Set hostname: vps-1051206.super.vz
VE start in progress...
[dagr@ssw2-ca ~]$ ps -ef| grep loop
root    25193    1 0 08:28 ?        00:00:00 [loop0]
root    25200    1 0 08:28 ?        00:00:00 [loop1]
root    27551    1 0 08:31 ?        00:00:00 [loop2]
root    27556    1 0 08:31 ?        00:00:00 [loop3]
dagr    29060 24812  0 08:31 pts/3    00:00:00 grep loop
[dagr@ssw2-ca ~]$ cat /etc/vz/conf/1051206.mount
#!/bin/bash
#echo mount
mount -n -o loop,ro /vz/shared_vps/php5_base/vps_static_loop_test /vz/root/$VEID/vps_static

mount -n -o loop,ro /vz/shared_vps/php5_base/vps_libs_loop /vz/root/$VEID/vps_libs
[dagr@ssw2-ca ~]$ cat /etc/vz/conf/1051206.umount
#!/bin/bash
echo umount
```

So eventually all 255 devices will be exhausted

Subject: Re: "vzctl stop" - hangs
Posted by [vaverin](#) on Thu, 15 Nov 2007 08:47:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

you can mount loop on the HW node and then bind-mount it to VEs via mount scripts

Subject: Re: "vzctl stop" - hangs
Posted by [dagr](#) on Thu, 15 Nov 2007 10:34:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hmm , thats a good idea anyway , but not sufficient in my case . One of that loopback mounts holds /usr,/sbin,/bin,/lib. So /etc/init.d/halt script just stalls because fuser shows him that fs still busy. So i had to change this script adding "halt; exit 0;" before umount of loopback devices.

Actually it helps even without using bind mounts. Just need to use mounts which do change /etc/mtab and in umounts add umount string for every loopback.

Anyway my problems are solved , tks.
