
Subject: VE refuses top stop, running init :: details included

Posted by [rickb](#) on Tue, 13 Nov 2007 16:53:38 GMT

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

Hi, I have seen this problem with the later 2.6.18 releases. After stopping a VE from HN vzctl interface or rebooting from within the VE, the init process refuses to exit. The only option is to kill the init process from the HN.

Quote:[root@berry ~]# vzlist 9891768

VEID	NPROC	STATUS	IP_ADDR	HOSTNAME
9891768	1	running	208.77.99.113	hostname.com

[root@berry ~]# lsof 2>/dev/null | grep 9891768

```
init  29566  root cwd    DIR    0,122   4096  49857524 /vz/root/9891768
init  29566  root rtd    DIR    0,122   4096  49857524 /vz/root/9891768
init  29566  root txt    REG    0,122  31216  50352681 /vz/root/9891768/sbin/init
init  29566  root mem    REG     9,2    50352681 /vz/root/9891768/sbin/init (path
dev=0,122)
init  29566  root mem    REG     9,2    52892116
/vz/root/9891768/lib/tls/libc-2.3.4.so (path dev=0,122)
init  29566  root mem    REG     9,2    50332266 /vz/root/9891768/lib/libselinux.so.1
(path dev=0,122)
init  29566  root mem    REG     9,2    50332292 /vz/root/9891768/lib/libsepol.so.1
(path dev=0,122)
init  29566  root mem    REG     9,2    50332296 /vz/root/9891768/lib/ld-2.3.4.so
(path dev=0,122)
init  29566  root 10u    FIFO    0,122    50383645 /vz/root/9891768/dev/initctl
```

[root@berry ~]# strace -ff -p 29566

Process 29566 attached - interrupt to quit

```
select(11, [10], NULL, NULL, {4, 983000}) = 0 (Timeout)
time(NULL)                                = 1194972738
stat64("/dev/initctl", {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
fstat64(10, {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
stat64("/dev/initctl", {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
select(11, [10], NULL, NULL, {5, 0})     = 0 (Timeout)
time(NULL)                                = 1194972743
stat64("/dev/initctl", {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
fstat64(10, {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
stat64("/dev/initctl", {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
select(11, [10], NULL, NULL, {5, 0})     = 0 (Timeout)
time(NULL)                                = 1194972748
stat64("/dev/initctl", {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
fstat64(10, {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
stat64("/dev/initctl", {st_mode=S_IFIFO|0600, st_size=0, ...}) = 0
select(11, [10], NULL, NULL, {5, 0}) <unfinished ...>
Process 29566 detached
```

```
[root@berry ~]#
```

```
[root@berry ~]# kill 29566
```

```
[root@berry ~]# ps ax | grep 29566
```

```
4629 pts/1 D+ 0:00 grep 29566
```

```
29566 ? Ss 0:06 init [3]
```

```
[root@berry ~]# kill -9 29566
```

```
[root@berry ~]# ps ax | grep 29566
```

```
4856 pts/1 S+ 0:00 grep 29566
```

```
[root@berry ~]# vzlist 9891768
```

```
    VEID    NPROC STATUS IP_ADDR    HOSTNAME  
9891768      - mounted 208.77.99.113 hostname.com
```

```
[root@berry ~]#
```

using 2.6.18 el5 release

any idears?

Rick

Subject: Re: VE refuses top stop, running init :: details included

Posted by [curx](#) on Tue, 13 Nov 2007 23:54:47 GMT

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

Hi Rick,

VE OS Template ?

Subject: Re: VE refuses top stop, running init :: details included

Posted by [rickb](#) on Wed, 14 Nov 2007 07:44:24 GMT

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

Hi, I know it happens on centos4, thats mostly what I run. I'll start to check each time it happens, I don't think its isolated to a VE distro.

Subject: Re: VE refuses top stop, running init :: details included

Posted by [vaverin](#) on Wed, 14 Nov 2007 08:03:19 GMT

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

Hi Rick,

"vzctl stop" should send SIGTERM and then SIGKILL.
But it isn't help if you make reboot inside VE.

I'll ask my colleguas about this case.

thank you,
Vasily Averin

Subject: Re: VE refuses top stop, running init :: details included
Posted by [vaverin](#) on Wed, 14 Nov 2007 08:04:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

kernel version is interesting too.

Subject: Re: VE refuses top stop, running init :: details included
Posted by [vaverin](#) on Wed, 14 Nov 2007 08:45:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

Some time ago we have disabled ability to kill init inside VE (like on standalone node where init ignores all signals). Probably this ability was used inside your VE. In this case you need to change initscripts and add some command that calls sys_reboot syscall (i.e reboot or halt).
