
Subject: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [rickb](#) on Tue, 12 Aug 2008 12:59:48 GMT

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

Hello, I am moving some VEs to another node the latest 64bit el5 patch set (moving from 32bit el5 openvz). When I start the VE on the new server, the VE fails to start from vzctl with error:

VE is mounted

VE start failed

VE is unmounted

In dmesg, the following error appears:

CT: 124069: stopped

CT: 124069: failed to start with err=-11

I have moved about 40 VEs from a 32bit host to this 64bit host, about 60% worked fine, the other 40% exhibit this error. The VEs are 32bit software and the OS templates range from ce5, ubuntu8, ce4, debian4, fedoras.

vzquota-3.0.11-1

vzctl-lib-3.0.22-1

vzctl-3.0.22-1

Any information is appreciated.

Rick

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [khorenko](#) on Tue, 12 Aug 2008 13:44:56 GMT

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

Hi Rick,

try to check failcounters and probably increase ulimits on HN.

--

Konstantin

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [rickb](#) on Tue, 12 Aug 2008 14:28:36 GMT

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

Thanks for the response.

The VEID is not running and is thus not in /proc/user_beancounters (no way to check failcnt).

```
[root@kool ~]# grep 124069 /proc/user_beancounters
[root@kool ~]#
```

ulimit on HN is unlimited.

```
[root@kool ~]# ulimit
unlimited
```

what triggers this error exactly?

Rick

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.
Posted by [khorenko](#) on Tue, 12 Aug 2008 14:42:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you please check/post `ulimit -a`?
Thank you.

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.
Posted by [rickb](#) on Tue, 12 Aug 2008 15:07:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thank you for the help.

```
[root@kool ~]# ulimit -a
core file size          (blocks, -c) 0
data seg size           (kbytes, -d) unlimited
scheduling priority     (-e) 0
file size               (blocks, -f) unlimited
pending signals         (-i) 137216
max locked memory       (kbytes, -l) 32
max memory size         (kbytes, -m) unlimited
open files              (-n) 1024
pipe size               (512 bytes, -p) 8
POSIX message queues    (bytes, -q) 819200
real-time priority      (-r) 0
stack size              (kbytes, -s) 10240
cpu time                (seconds, -t) unlimited
max user processes      (-u) 137216
```

virtual memory (kbytes, -v) unlimited
file locks (-x) unlimited
[root@kool ~]#

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.
Posted by [rickb](#) on Tue, 12 Aug 2008 15:36:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi, I tried to create a new VE and I have the same error. There are exactly 46 running VEs.

if I stop a running VE, I can start my new one. But I cannot start the one I just stopped.

So, it seems there is some limit imposed at exactly 46 running VE.

Rick

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.
Posted by [rickb](#) on Tue, 12 Aug 2008 15:43:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

Seems this error is due to a kcalloc() failure. should I downgrade the kernel to the last release and retry?

high mem reports 0/0..?

```
[root@kool src]# cat /proc/meminfo
MemTotal: 16405428 kB
MemFree: 5352608 kB
Buffers: 1575036 kB
Cached: 5155936 kB
SwapCached: 1596 kB
Active: 6639644 kB
Inactive: 3217932 kB
HighTotal: 0 kB <<=====
HighFree: 0 kB <<=====
LowTotal: 16405428 kB
LowFree: 5352608 kB
SwapTotal: 4096524 kB
SwapFree: 4081160 kB
Dirty: 5712 kB
Writeback: 552 kB
AnonPages: 3124744 kB
Mapped: 579088 kB
```

Slab: 1008840 kB
PageTables: 56512 kB
NFS_Unstable: 0 kB
Bounce: 0 kB
CommitLimit: 12299236 kB
Committed_AS: 12543696 kB
VmallocTotal: 34359738364 kB
VmallocUsed: 308236 kB
VmallocChunk: 34359425572 kB
HugePages_Total: 0
HugePages_Free: 0
HugePages_Rsvd: 0
Hugepagesize: 2048 kB

Rick

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [khorenko](#) on Tue, 12 Aug 2008 20:14:51 GMT

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

Hi again Rick,

you are running x86_64 kernel and all available memory on the node can be accessed directly as lowmem as there is no 4Gb addressing limitation. Thus highmem is reported as 0. That's normal.

Unable to start more than 46 Containers... this is definitely some resource limitation. Can you please try just in case to increase all your ulimits in two times and try once more? Just to be sure that the problem is not here.

Could you please also get a strace of vzctl start for the same Container when it goes ok and when it fails?

If this info won't help us i'll probably create a debug kernel for you to track which exactly place fails in kernel.

Thank you!

--

Konstantin

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [rickb](#) on Tue, 12 Aug 2008 20:26:38 GMT

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

Hi, I doubled all ulimit values, I see the same problem.

```
[root@kool bc]# ulimit -a
core file size          (blocks, -c) 0
data seg size           (kbytes, -d) unlimited
scheduling priority      (-e) 0
file size                (blocks, -f) unlimited
pending signals          (-i) 237216
max locked memory        (kbytes, -l) 64
max memory size          (kbytes, -m) unlimited
open files               (-n) 2048
pipe size                (512 bytes, -p) 8
POSIX message queues     (bytes, -q) 1619200
real-time priority       (-r) 0
stack size               (kbytes, -s) 20480
cpu time                 (seconds, -t) unlimited
max user processes       (-u) 237216
virtual memory           (kbytes, -v) unlimited
file locks                (-x) unlimited
```

```
[root@kool bc]# vzctl start 9893392 --force
Starting VE ...
VE is mounted
VE start failed
VE is unmounted
[root@kool bc]#
[root@kool bc]# dmesg | grep 9893392 | tail -10
CT: 9893392: stopped
CT: 9893392: failed to start with err=-11
CT: 9893392: stopped
CT: 9893392: failed to start with err=-11
CT: 9893392: stopped
CT: 9893392: failed to start with err=-11
[root@kool bc]#
```

I am going to downgrade the kernel tonight as I found another bug in this one, the read/write accounting in `/proc/bc/<veid>/ioacct` is nonexistent.

```
[root@kool bc]# egrep 'read|write' */ioacct | grep vfs -v | head -10
0/ioacct:   read          0
0/ioacct:   write         0
402099/ioacct: read       0
402099/ioacct: write      0
6813/ioacct: read        0
6813/ioacct: write       0
889205/ioacct: read       0
889205/ioacct: write      0
889431/ioacct: read       0
```

This is with patch patch-92.1.1.el5.028stab057.2-combined, I am going to downgrade to patch-53.1.19.el5.028stab053.14-combined tonight and I will report back if either of these problems are solved by this.

Rick

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [khorenko](#) on Tue, 12 Aug 2008 21:08:47 GMT

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

-
- 1) Try vzctl --verbose start ?
 - 2) straces of 'vzctl start'?
 - 3) btw, don't you have some quota enabled on your HN? For example quota for number of users on a system?

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [rickb](#) on Tue, 12 Aug 2008 21:22:15 GMT

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

finist, thanks for your help. I will post 1&2 soon. #3 is no.

Rick

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [maratrus](#) on Wed, 13 Aug 2008 06:24:40 GMT

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

Hello,

how did you manage do find out that the problem is in the kcalloc()?
May be it's worth increasing printk level and vz log level too (/etc/vz/vz.conf)?

Subject: Re: VE Fails to start with dmesg error "err=-11". 64bit el5 patch.

Posted by [rickb](#) on Wed, 13 Aug 2008 08:22:32 GMT

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

Hi, seems like the code is failing kcalloc here:

Quote: err = -ENOMEM;
 ve = kzalloc(sizeof(struct ve_struct), GFP_KERNEL);
 if (ve == NULL)
 goto err_struct;
and jumping to:

Quote:err_struct:
 printk(KERN_INFO "CT: %d: failer to start with err=%d\n", veid, err);
 return err;

strace output from the failing "vzctl start <veid>" command is below:
<http://208.77.101.170/str.txt>

I booted off el5 patched with patch-53.1.19.el5.028stab053.14-combined and the system does not exhibit either problem I reported (ioacct accounting or velimit). Due to this I will not be able to test any patches, but I am working on getting a development 64bit node up to test should you release one.

Rick
