Subject: *KNOWN ISSUE* Failed to dump VE on 028stab033 kernel Posted by goeldi on Fri, 01 Jun 2007 15:15:19 GMT

View Forum Message <> Reply to Message

I tried again to vzmigrate online between 2 hn:

On both HN runs 2.6.18-ovz028stab033.1-smp

This is what happened:

nice vzmigrate -r no --online -v 195.141.118.235 343

OPT:-r

OPT:--online

OPT:-v

OPT:195.141.118.235

Starting online migration of VE 343 on 195.141.118.235

OpenVZ is running...

Loading /etc/vz/vz.conf and /etc/vz/conf/343.conf files

Check IPs on destination node: 195.141.143.43

Preparing remote node

Copying config file

343.conf 100% 1378 1.4KB/s 00:00

Saved parameters for VE 343

Creating remote VE root dir

Creating remote VE private dir

Initializing remote quota

Quota init

vzquota: (warning) Quota file exists, it will be overwritten

Turning remote quota on

Syncing private

Live migrating VE

Suspending VE

Setting up checkpoint...

suspend...

get context...

Checkpointing completed succesfully

Dumping VE

Setting up checkpoint...

join context..

dump...

Can not dump VE: Invalid argument

iptables-save exited with 16180

Checkpointing failed

Error: Failed to dump VE

Resuming...

Subject: Re: Failed to dump VE

Posted by Vasily Tarasov on Mon, 04 Jun 2007 06:33:40 GMT

View Forum Message <> Reply to Message

Hello,

It's a known issue: on 028stab033 checkpointing is broken! It's fixed in git already with the following patch:

http://git.openvz.org/?p=linux-2.6.18-openvz;a=commit;h=650723f83a8685111fb3dee7073454861c446193

So, in order to use checkpointing you can:

- 1) use 031 kernel before 034 is released
- 2) compile fresh kernel from git

Sorry for inconvenience!

Vasily.

Subject: Re: Failed to dump VE

Posted by goeldi on Mon, 04 Jun 2007 07:06:53 GMT

View Forum Message <> Reply to Message

Thank you. Unfortunately with 031 I have a massive I/O problem

Subject: Re: Failed to dump VE

Posted by Vasily Tarasov on Mon, 04 Jun 2007 07:09:32 GMT

View Forum Message <> Reply to Message

Yes, it's really a bad situation... Sorry.

Subject: Re: Failed to dump VE

Posted by hernad on Wed, 06 Jun 2007 09:11:16 GMT

View Forum Message <> Reply to Message

I have built the newest kernel from git to fix this issue.

I am using ubuntu feisty, making packages with:

~/openvz/linux-2.6.18-openvz-033-b\$ make-kpkg -initrd linux_image linux_headers --rootcmd fakeroot

Trying to boot new kernel, I'm getting boot error:

init_premount / udev 21 cannot create /proc/sys/kernel/hotplug ... operation not permited

After some time, I am getting boot prompt (initramfs)

and message:

cannot find: /dev/mapper/Ubuntu-root

System obviously cannot find my SATA disk and LVM volumes.

I have used the same procedure and same .config to build kernel packages stable033 and stable031, and I had no problems.

Regards, Ernad

Subject: Re: Failed to dump VE

Posted by Vasily Tarasov on Wed, 06 Jun 2007 09:54:45 GMT

View Forum Message <> Reply to Message

Can you detect who and when does produce the message "init_premount / udev 21 cannot create /proc/sys/kernel/hotplug ... operation not permitted" Is it initrd? Or init already?

I can say you it if you post here the full log of kernel boot...

Thank you, Vasily.

Subject: Re: Failed to dump VE Posted by hernad on Wed, 06 Jun 2007 10:48:47 GMT

View Forum Message <> Reply to Message

It is initrd script.

Here is the script:

hernad@desk-h8:~/openvz/temp\$ gzip -dc /boot/initrd.img-2.6.18-033-b | cpio -id 22897 blocks

hernad@desk-h8:~/openvz/temp\$ cat scripts/init-premount/udev #!/bin/sh -e # initramfs premount script for udev

PREREQ=""

```
# Output pre-requisites
prereqs()
    echo "$PREREQ"
}
case "$1" in
  prereqs)
     preregs
     exit 0
esac
# It's all over netlink now
echo "" > /proc/sys/kernel/hotplug <<<< LINE 21
# Start the udev daemon to process events
/sbin/udevd --daemon
# Iterate sysfs and fire off everything; if we include a rule for it then
# it'll get handled; otherwise it'll get handled later when we do this again
# in the main boot sequence.
/sbin/udevtrigger
I have checked, thre is no difference between initrd.img of this 033-HEAD, stab033 i stab031
kernel.
I don't know how to provide you log of initrd process.
System cannot acces root file system (/dev/mapper/Ubuntu-root).
Thank you.
Subject: Re: /proc entries can't be modified! (was: Failed to dump VE)
Posted by curx on Wed, 06 Jun 2007 10:54:50 GMT
View Forum Message <> Reply to Message
Hi,
it's a know issue, please check:
```

http://bugzilla.openvz.org/show_bug.cgi?id=604

Subject: Re: Failed to dump VE

Posted by Vasily Tarasov on Wed, 06 Jun 2007 10:59:18 GMT

View Forum Message <> Reply to Message

Ouch!... Now I get the problem. The following patch from git:

http://git.openvz.org/?p=linux-2.6.18-openvz;a=commitdiff;h=

438dfab37c427b0ec6e35cc4898f240c84367c57

breaks the kernel: nobody can't write to /proc! It's already fixed:

http://bugzilla.openvz.org/show_bug.cgi?id=604. Please, get the patch from the bug report and compile the kernel with it.

Sorry for inconvenience, Vasily

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE) Posted by hernad on Wed, 06 Jun 2007 11:06:03 GMT

View Forum Message <> Reply to Message

Hm, It seems this patch is already applied to my source ?!

hernad@desk-h8:~/openvz/linux-2.6.18-openvz-033-b\$ patch -p1 < 1.patch patching file fs/proc/generic.c Reversed (or previously applied) patch detected! Assume -R? [n]

I am using kernel master branch from git, last updated this morning:

hernad@desk-h8:~/openvz/git/linux-2.6.18-openvz\$ git show-branch [master] syslog: fix LTP failure in case of wrong args

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE) Posted by Vasily Tarasov on Wed, 06 Jun 2007 11:12:51 GMT

View Forum Message <> Reply to Message

I just attached to the bug report the right patch, that should be appliable to your kernel.

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE)

Posted by hernad on Thu, 07 Jun 2007 08:20:38 GMT

View Forum Message <> Reply to Message

Thank you Vasily,

with this patch, boot problem is resolved.

But checkpointing still doesn't work ?!

root@desk-h8:~# vzctl chkpnt 999 --dump --dumpfile /vz/999_dump Setting up checkpoint...

join context..

dump...

Can not dump VE: Invalid argument

iptables-save exited with 255

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE) Posted by hernad on Thu, 07 Jun 2007 08:21:50 GMT

View Forum Message <> Reply to Message

I reverted to stab031 kernel, and migration finally works.

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE) Posted by Vasily Tarasov on Thu, 07 Jun 2007 08:48:05 GMT

View Forum Message <> Reply to Message

Ahhhh... Sorry, I noted that current git doesn't content required patch... Please, wait till 034 kernel, it'll be very soon.

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE)

Posted by hernad on Thu, 07 Jun 2007 09:10:40 GMT

View Forum Message <> Reply to Message

No problem

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE)

Posted by hernad on Fri, 08 Jun 2007 15:34:08 GMT

View Forum Message <> Reply to Message

I have not success with live migration with stab035 from git.

vzctl reported the sam problem as with stab033.

Maybe I should get the new vzctl?

Regards, Ernad

File Attachments

1) error_stab035.txt, downloaded 435 times

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE) Posted by hernad on Fri, 15 Jun 2007 08:36:49 GMT

View Forum Message <> Reply to Message

2.6.18-stab036 kernel, problem remains

Can not dump VE: Invalid argument iptables-save exited with 255

Subject: Re: /proc entries can't be modified! (was: Failed to dump VE) Posted by hernad on Fri, 15 Jun 2007 08:45:06 GMT

View Forum Message <> Reply to Message

http://bugzilla.openvz.org/show_bug.cgi?id=621