Subject: Even worse thing when migrating online
Posted by divB on Sat, 09 May 2009 22:04:42 GMT
View Forum Message <> Reply to Message

Hi,

Now I have a problem even much more worse (does this exist?). Sometimes when I do online-migration (from HN1 to HN2) the network connection between the two hosts drops (the funny thing: ONLY between the two hardware nodes!) and this leads to a fatal situation:

* The ssh commands from vzmigrate are not executed any more
* The VE is still up on HN1
* But it is also up on HN2 as "zombie VE"

```
HN1:~# vzlist
      VEID      NPROC STATUS  IP_ADDR         HOSTNAME
       201          6 running -


HN2:~# vzlist
      VEID      NPROC STATUS  IP_ADDR         HOSTNAME
       201          9 running -
HN2:~# vzctl enter 201
enter into VE 201 failed
HN2:~#
```

This is where the migration looks like:

```
NH1:~# vzmigrate2 -r no --keep-dst --online -v 192.168.200.1 201
OPT:-r
OPT:--keep-dst
OPT:--online
OPT:-v
OPT:192.168.200.1
Starting online migration of VE 201 on 192.168.200.1
OpenVZ is running...
   Loading /etc/vz/vz.conf and /etc/vz/conf/201.conf files
   Check IPs on destination node:
Preparing remote node
   Copying config file
201.conf                                                      100% 1756
1.7KB/s   00:00
Saved parameters for VE 201
   Creating remote VE root dir
   Creating remote VE private dir
```

```
   VZ disk quota disabled -- skipping quota migration
Syncing private
Live migrating VE
Stop apache2 if it is installed
Stopping web server: apache2 ... waiting .
   Suspending VE
Setting up checkpoint...
        suspend...
        get context...
Checkpointing completed succesfully
   Dumping VE
Setting up checkpoint...
        join context..
        dump...
Checkpointing completed succesfully
   Copying dumpfile
dump.201                                                      100% 1492KB
1.5MB/s   00:01
   Syncing private (2nd pass)
   VZ disk quota disabled -- skipping quota migration
   Undumping VE
Restoring VE ...
Starting VE ...
VE is mounted
        undump...
Setting CPU units: 1000
Configure meminfo: 2147483647
Configure veth devices: veth201.0
        get context...
VE start in progress...
Restoring completed succesfully
Adding interface veth201.0 to bridge br-lan on CT0 for CT201
```

After that, the script hangs. Clearly, as said, pinging HN2 is not possible any more. This leads to a hang of the SSH commands:

```
HN1:~# ps aux
[...]
root      3914  0.2  0.1   3928  1320 pts/1    S+   01:43   0:00 /bin/sh /usr/local/sbin/vzmigrate2 -r no
--keep-dst --online -v 192.168.200.1 201
root      3974  0.2  0.2   5124  2288 pts/1    S+   01:43   0:00 ssh root@192.168.200.1 vzctl restore
201 --undump --dumpfile /var/tmp/dump.201 --skip_arpdet
```

After killing PID 3974, the next ssh command from the vzmigrate script is spawned:

```
HN1:~# ps aux
[...]
root     3914  0.1  0.1   3928  1320 pts/1    S+   01:43   0:00 /bin/sh /usr/local/sbin/vzmigrate2 -r no
--keep-dst --online -v 192.168.200.1 201
root     3975  0.0  0.1   4248  1676 pts/2    Ss   01:43   0:00 /bin/bash
root     3978  6.0  0.1   5124  1828 pts/1    S+   01:44   0:00 ssh root@192.168.200.1 rm -f
/var/tmp/quotadump.201
```

As mentioned above, both hardware nodes are now inconsistent and "buggy". Just deleting /etc/vz/conf/201.conf and then rebooting BOTH hardware nodes resolves the problem

Well, but what exactly happens when starting my machines? First I have to mention that I only use vzeth and no vznet. So I have to make sure to bridge the veth-Device together with the bridges on the hardware node.

Additionally I have to big problem that Debian lenny does not yet support the EXTERNAL_SCRIPT functionality. So I hacked the wurgaround I found in [1].

So in common, my /etc/vz/conf/vps.mount looks like [2].

In this script, the vznetaddbr explained in [1] is called. The contents of this file is in [3].

The very big question now: Why does this happen? From a third computer I can ping both hardware nodes but they can't communicate anymore with each other! I am not sure if this problem is caused my bridging scripts...

Is there any hope to resolve this issue?

Thank you very much,
divB

[1] http://wiki.openvz.org/Veth#method_for_vzctl_version_.3C.3D_ 3.0.22
[2] http://pastebin.com/m33a4232a
[3] http://pastebin.com/m2136da98