Subject: NFS Outside ve
Posted by ready on Mon, 15 Jan 2007 15:11:17 GMT
View Forum Message <> Reply to Message

Hi,

i'm planning to mount the whole vz directory through nfs.
The idea behind the scenario is having the vz directory stored on a raid5 nas server and several boxes running vz kernel which can run a ve from nas vz tree.
So i can easily swich one ve to another box if the load on the box is too heavy.

I did that by the following way:
I have set up a box with the vz stuff and tested the ve's.
Then i rsynced the /var/lib/vz directory to the nas server.
I exported /srv/vz by:
/srv/vz box*(rw,sync,no_root_squash,no_all_squash,subtree_check)
I mounted the /srv/vz to a vz box by:
mount nas1:/srv/vz /var/lib/vz

And started one ve. Now, i get some really strange behavoir:

box1:~# vzctl start 500
Removing stale loroot@auth:/#ck file /var/lib/vz/lock/500.lck
Starting VE ...
VE is mounted
Adding IP address(es): 192.168.66.199
Setting CPU units: 1000
Set hostname: auth
File resolv.conf was modified
VE start in progress...

box1:~# vzctl enter 500
entered into VE 500
root@auth:/# l
total 80
drwxr-xr-x 2 root root 4096 Dec 18 08:27 bin
drwxr-xr-x 2 root root   6 Dec 15 2004 boot
drwxr-xr-x 11 root root 24576 Jan 15 2007 dev
drwxr-xr-x 47 root root 4096 Jan 15 2007 etc
drwxrwsr-x 2 root staff   6 Dec 15 2004 home
drwxr-xr-x 2 root root   6 Dec 20 2005 initrd
drwxr-xr-x 9 root root 4096 Nov 24 11:40 lib
drwxr-xr-x 2 root root   6 Dec 20 2005 media
drwxr-xr-x 2 root root   6 Dec 15 2004 mnt
drwxr-xr-x 2 root root   6 Dec 20 2005 opt
dr-xr-xr-x 76 root root   0 Jan 15 15:37 proc
drwx------ 3 root root   111 Jan 2 17:21 root
drwxr-xr-x 2 root root 4096 Dec 18 08:27 sbin

```
drwxr-xr-x  2 root root      6 Dec 20  2005 srv
drwxr-xr-x 11 root root      0 Jan 15 15:37 sys
drwxrwxrwt  2 root root      6 Jan 15  2007 tmp
drwxr-xr-x 11 root root    107 Nov 16 08:54 usr
drwxr-xr-x 14 root root    138 Nov 17 12:19 var
root@auth:/# logout
exited from VE 500
```

i got logged out automatically after i listed the directory.
It doesnt matter what i type, i always get logged out after the first command when i enter the fresh started ve.

Well, i reenter:

```
box1:~# vzctl enter 500
entered into VE 500
root@auth:/# ps a
  PID TTY      STAT   TIME COMMAND
 4880 pts/0    Ss     0:00 -bash
 4984 tty2     Ss+    0:00 /sbin/getty 38400 tty2
 4985 tty3     Ss+    0:00 /sbin/getty 38400 tty3
 4986 tty4     Ss+    0:00 /sbin/getty 38400 tty4
 4987 tty5     Ss+    0:00 /sbin/getty 38400 tty5
 4988 tty6     Ss+    0:00 /sbin/getty 38400 tty6
 6724 tty1     Ss+    0:00 /sbin/getty 38400 tty1
14730 pts/1    Ss     0:00 -bash
14734 pts/1    S+     0:00 less /var/log/syslog
14778 pts/0    R+     0:00 vzctl enter 500
15804 pts/0    Ss     0:00 -bash
root@auth:/#
```

The really strange thing is that i get the whole process tree, now only the ve processes!!

I leave the ve:

```
root@auth:/# logout
exited from VE 500
box1:~#
```

And i want to shut down the VE:

```
box1:~# vzctl stop 500
Stopping VE ...
```

Connection to box1 closed by remote host.
Connection to box1 closed.

Now the whole System hangs up!

Are the problems known and are they going to fixed?
And if not, which alternatives there are to get the above described scenario?

Some infos:
vz kernel: 2.6.18-028test010.1-ovz-smp
system: Debian etch
Filesystem: xfs

Any help appreciated

Greets

Subject: Re: NFS Outside ve
Posted by dca@adviseo.fr on Tue, 06 Feb 2007 16:02:53 GMT
View Forum Message <> Reply to Message

Hi,

I get the same error:

No problems if i use an XFS partition to store VE privates directories, but when accessed mounted localy with NFS, VEs use the hosts's root and not their private one. IP configuration also comes on the host... and everything crash.

Any idea ?

kernel 2.6.18-028test010
debian etch
XFS
nfs-kernel server into host

Thanks

Subject: Re: NFS Outside ve
Posted by Vasily Tarasov on Wed, 07 Feb 2007 07:33:04 GMT
View Forum Message <> Reply to Message

Hello,

Thank you all for this report. I filled the bug http://bugzilla.openvz.org/show_bug.cgi?id=464. Add you to CC if you want to be aware about the progress on the bug.

Thanks again,
Vasily
