
Subject: ovzkernel 2.6.16-026test009.1 and RHEL 4 update 3 problem

Posted by [dowdle](#) on Thu, 20 Apr 2006 19:51:49 GMT

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

Greetings,

I installed ovzkernel 2.6.16-026test009.1 on two machines today. One was running CentOS 4.3 and the other Red Hat Enterprise Linux 4 Update 3... completely different hardware.

Anyway...

Here's the problem I'm having on the RHEL machine.

- 1) Apache is in a VPS
- 2) The physical host has several NFS mounts
- 3) The NFS mounts are bound to the VPS
- 4) Apache in the VPS has one of the NFS mounts as DocumentRoot
- 5) The web pages are serving the text portions just fine but embedded .jpg's in pages just sit and spin... and eventually time out. I can reload the page and get some of the images to show up.
- 6) Nothing looks out of the ordinary in /var/log/httpd/*
- 7) Nothing looks out of the ordinary in the physical host logs
- 8] /proc/user_beancounters is all 0's

On the first attempt to install the test kernel, I got an rpm conflict with a package named:

lksctp-tools

I removed that package and the test kernel installed just fine.

Any idea what is going on? I've restarted apache and the VPS but that didn't help. NFS is working fine... and the files are available in the VPS... apache just is not transferring them... or the HTTP protocol is getting confused or something.

All other services appear to be working fine so I'm assuming it isn't some sort of connectivity issue. iptables is NOT running on the physical host at this time.

Any ideas?

I did not get the package conflict when I installed the test kernel on CentOS 4.3 and oddly enough the RHEL install was a minimal install... so the package I had to remove must have been a default?!? Of course, like I said... the CentOS is completely different hardware and isn't doing the NFS stuff at all.

Subject: Re: ovzkernel 2.6.16-026test009.1 and RHEL 4 update 3 problem

Posted by [kir](#) on Thu, 20 Apr 2006 20:44:37 GMT

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

lksctp-tools is a set of utils to support SCTP protocol (a recently new protocol). You probably do not want it. If you do not know -- then you do not want it.

We will look into the conflict -- probably the package requires something from the kernel (SCTP support?), well, from the rpm package, which OpenVZ kernel package misses. Indeed, we do not support SCTP.

Speaking of apache problem, it requires a separate investigation. Probably it is NFS -- can you (temporary) try it with local files?

Subject: Local files, no problem: NFS files, problems
Posted by [dowdle](#) on Thu, 20 Apr 2006 22:03:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

I used wget and dragged some of the files over into a local directory and then modified Apache's DocumentRoot... restarted apache... and no problem at all.

I switch it back, and problems.

The same configuration did not give me a problem with the stable OpenVZ kernel.

Another question... sorta related but not.

Ok, in /etc/sysconfig/vz-scripts I have 101.mount and 101.umount and both are 755.

Check out the following behavior:

(before starting 101 - physical host)

```
[root@hprack6 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/cciss/c0d0p1 30G  1.8G  27G   7% /
none            1012M  0 1012M   0% /dev/shm
esus:/students  145G  31G  107G  23% /mnt/nfs/esus/students
esus:/home      145G  112G  26G  82% /mnt/nfs/esus/home
esus:/www/htdocs 55G  11G  43G  20% /mnt/nfs/esus/www
[root@hprack6 ~]#
```

(after starting vps - physical host)

```
[root@hprack6 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/cciss/c0d0p1 30G  1.8G  27G   7% /
none            1012M  0 1012M   0% /dev/shm
esus:/students  145G  31G  107G  23% /mnt/nfs/esus/students
esus:/home      145G  112G  26G  82% /mnt/nfs/esus/home
esus:/www/htdocs 55G  11G  43G  20% /mnt/nfs/esus/www
/mnt/nfs/esus/home 145G  112G  26G  82% /vz/root/101/home
/mnt/nfs/esus/students
```

```
145G 31G 107G 23% /vz/root/101/students
/mnt/nfs/esus/www 55G 11G 43G 20% /vz/root/101/www
[root@hprack6 ~]#
```

(after starting the vps - vps)

```
[root@web2 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
simfs           5.0G 426M 4.6G  9% /
nfs             145G 112G  26G 82% /home
nfs             145G  31G 107G 23% /students
nfs             55G  11G  43G 20% /www
[root@web2 ~]#
```

(after stopping vps - physical host)

```
[root@hprack6 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/cciss/c0d0p1 30G 1.8G 27G  7% /
none            1012M  0 1012M  0% /dev/shm
esus:/students  30G 1.8G 27G  7% /mnt/nfs/esus/students
esus:/home      30G 1.8G 27G  7% /mnt/nfs/esus/home
esus:/www/htdocs 30G 1.8G 27G  7% /mnt/nfs/esus/www
[root@hprack6 ~]#
```

Notice how all of info for the nfs mounts has changed and oddly enough they reflect the / filesystem rather than their real values? CDing into those dirs and doing an ls shows NO files... everything has disappeared. When I umount the dirs by hand they say they aren't mounted... yet they disappear from the df listing... and then I mount -a and they get mounted correctly again.

Here are the contents of my 101.mount script:

```
#!/bin/bash
#
# 101.mount - script to bind nfs mounts of physical host
# if one of these files does not exist then something is
# really broken
[ -f "/etc/sysconfig/vz-scripts/101.conf" ] || exit 1
# If home dirs are not mounted we exit with error
mount --bind /mnt/nfs/esus/home /vz/root/101/home
mount --bind /mnt/nfs/esus/students /vz/root/101/students
mount --bind /mnt/nfs/esus/www /vz/root/101/www
exit $?
```

Here are the contents of my 101.umount script:

```
#!/bin/bash
# if one of these files does not exist then something is really broken
[ -f "/etc/sysconfig/vz-scripts/101.conf" ] || exit 1
# umount nfs mounts
umount /vz/root/101/home
umount /vz/root/101/students
```

```
umount /vz/root/101/www
exit $?
```

Any idea what is going on? Sorry for the bad formatting. The <pre> tag doesn't seem to work for me here.

I'm not sure if this is related to the test kernel or if I had the same problem with the stable kernel. I installed the new vzctl with the migration features... which requires the test kernel. If I reboot and run the stable kernel am I in need of downgrading vzctl anv vzctl-libs?

Subject: Re: Local files, no problem: NFS files, problems
Posted by [kir](#) on Fri, 21 Apr 2006 06:55:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Looks like the problems with kernel, not vzctl. Can you file two bugs to bugzilla? First one is about apache on nfs -- please specify versions where it's working and where it's not. Second one is about nfs being bind-mounted and then unmounted.

PS here you can use [CODE] ... [/CODE] pair of tags to preformat the pieces of code/logs/output. Just remove the spaces from [CODE].

Subject: Re: Local files, no problem: NFS files, problems
Posted by [dowdle](#) on Fri, 21 Apr 2006 20:16:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

Filed bugs #143 and #144.

Subject: Re: ovzkernel 2.6.16-026test009.1 and RHEL 4 update 3 problem
Posted by [dev](#) on Sat, 22 Apr 2006 08:06:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Scott, I suppose we need to check 2 things:

1. SendFile. We need to check whether apache is using or not SendFile. I'm not a big apache guru, but AFAIK there should be SendFile option for this in httpd.conf. Can you check it please? If it uses SendFile try to switch it OFF. This will make us closer to the problem

BTW, do you use i686 installation and kernel? just in case...

2. there should be something written about this in apache log. If it fails to send data, then it should log. If it is stuck, then you can check whether there are processed in D state with ps. We need to find what apache does and where the problem occurs. You can also set MaxClients to 1 in apache configs, restart it and strace the only apache child.

Subject: Re: Local files, no problem: NFS files, problems

Posted by [dev](#) on Sat, 22 Apr 2006 08:09:32 GMT

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

you don't need downgrading vzctl etc. everything should be backward compatible.

Subject: Re: ovzkernel 2.6.16-026test009.1 and RHEL 4 update 3 problem

Posted by [dowdle](#) on Mon, 24 Apr 2006 22:25:22 GMT

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

Ok, I uncommented the following line in the apache config:

```
EnableSendfile off
```

I wasn't familiar with that option until you brought it to my attention. I restarted apache and that didn't seem to make a difference.

I checked the status of the running apache processes and none of them were in a D state... they were all sleeping.

I'm not that familiar with using strace... so I didn't get anything useful... and it seemed to detach almost immediately.

I tailed the access_log while browsing pages... sometimes the pages would show as 200 (sent) but often my browser didn't get them... sometimes the request was never received by the server.

I had a kernel Ooops and had to reboot. There was nothing intelligible in the oops message. I decided to boot to the stable kernel.

I did double check that apache in a vps has no problem serving up the files from the stable kernel... but oddly enough the weird behavior with NFS mounts *IS* there with the stable kernel.

- 1) on physical host, mounts look fine
- 2) stop my vps
- 3) all of the mounts on the physical host reflect the info of the / filesystem
- 4) I have to umount all nfs mounts and mount them again for them to work again

Weird.
