

---

Subject: simfs problems with 3.7 TiB NFS share  
Posted by [Sten](#) on Fri, 14 Dec 2007 00:03:26 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I have two servers, one disk server (DS) with 3.7 TiB NFS share and one application server (AS) running OpenVZ. I have multiple VEs on the AS running from the NFS and some running from local disk.

The VEs running from ASs disks (2 GiB RAID 1) works without any problem, however the VEs running from the NFS share has problems with displaying disk size; df fails on 'Unknown error 521':

```
Thufir / # LANG=C df
Filesystem      1K-blocks    Used Available Use% Mounted on
tmpfs           4076832      0 4076832  0% /lib/init/rw
tmpfs           4076832      0 4076832  0% /dev/shm
df: `/' : Unknown error 521
```

The same problem occurs even on the AS (the host) itself; note that the NFS itself works fine:

```
Mentat samba # LANG=C df -h
Filesystem      Size  Used Avail Use% Mounted on
udev            3.9G  48K  3.9G   1% /dev
/dev/mapper/Mentat 1.9G  1.1G  685M  62% /union/ro
none            3.9G  17M  3.9G   1% /union/rw
union           3.9G  17M  3.9G   1% /
tmpfs           3.9G   0  3.9G   0% /lib/init/rw
tmpfs           3.9G  4.0K  3.9G   1% /dev/shm
/dev/md0        69M   9.5M  56M  15% /boot
/dev/md0        69M   9.5M  56M  15% /union/ro/boot
udev            3.9G  48K  3.9G   1% /union/ro/dev
/var/lib/vz/private/22002
                 3.9G  17M  3.9G   1% /var/lib/vz/root/22002
/dev/mapper/Mentat 1.9G  1.1G  685M  62% /var/lib/vz/root/22002/ro
tmpfs           3.9G   0  3.9G   0% /var/lib/vz/root/22002/lib/init/rw
tmpfs           3.9G   0  3.9G   0% /var/lib/vz/root/22002/dev/shm
192.168.193.2:/pool/Miles
                 3.7T 484G  3.2T  13% /pool/Miles
df: `/var/lib/vz/root/74002': Unknown error 521
tmpfs           3.9G   0  3.9G   0% /var/lib/vz/root/74002/lib/init/rw
tmpfs           3.9G   0  3.9G   0% /var/lib/vz/root/74002/dev/shm
192.168.193.2:/pool/Hayt
                 3.7T 484G  3.2T  13% /pool/Hayt
df: `/var/lib/vz/root/74003': Unknown error 521
```

and many similar errors differing only in VEID and server name

There are no information in dmesg or syslog regarding this. The OpenVZ kernel I'm running is

2.6.22-ovz005 with MadWifi and UnionFS patches. The kernel on the DS is 2.6.22.14 vanilla.

UPDATE: Both servers are dual core AMD64 with 64bit Debian. All VEs are 64bit Debian as well.

/etc/exports on the DS

/pool 192.168.193.1(rw, sync, no\_subtree\_check, no\_root\_squash)

The VEs has no quota at all.

ANOTHER UPDATE: When the VE is run from smaller (local) drive and the NFS is mounted inside it, it works fine. I haven't tested whether the problem remains when simfsing such disk array directly.

If the problem would affect only the df, it would be inconvenient but one could live with it, but it affects all applications that wants to know how big the share is. So the big deal that's killing me is that Samba running in one such VE permanently crashes when anyone tries to open any file in RW mode (and that it does not report the disk size at all).

So the question is:

is there any possibility to fix some configuration and make simfs happy? or is this a bug somewhere in something?

---

Subject: Re: simfs problems with 3.7 TiB NFS share  
Posted by [vaverin](#) on Fri, 14 Dec 2007 09:24:46 GMT  
[View Forum Message](#) <> [Reply to Message](#)

include/linux/errno.h:

/\* Defined for the NFSv3 protocol \*/

#define EBADHANDLE 521 /\* Illegal NFS file handle \*/

However currently I'm not sure that I understand how do You configured VE over NFS. (you mounts NFS share to the host system and then use its content as VE's private area, is I understand You correctly?)

Could You please create new bug in our bugzilla ([bugzilla.openvz.org](http://bugzilla.openvz.org)) and attach strace df, your kernel config, VE configs and mount scripts.

thank You,  
Vasily Averin

---

Subject: Re: simfs problems with 3.7 TiB NFS share  
Posted by [Sten](#) on Fri, 14 Dec 2007 13:51:33 GMT  
[View Forum Message](#) <> [Reply to Message](#)

Oh, strange that it's not defined in glibc.

The VEs are configured by a script of mine that mounts a NFS share and than starts a VE inside it. So, the VE is mapped to the mount point (but it does not work when the VE is mapped to a subdirectory of such mount point either).

I'll send a bug report the moment I'll get to a place where I can obtain SSH access to the servers.