
Subject: Mounts--bind, issue or limitation?

Posted by [3molo](#) on Sun, 06 May 2007 22:11:57 GMT

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

Hey.

Got 2x 500GB IDE disks, using dmccrypted (sha256) ext3 filesystem.

Mounted as /data1 & /data2 on the HN, the mount --bind into /vz/root/VEID/. After a short period of time, the owners of the VPS in which I have mount bind into started

to do alot of mount binds because glftpd has a hard time with symlinks(or something).

This resulted in me not being able to use the disks at all on the hardware node. Output of mount and df below:

Since I rather not interfer unless I have to, I wonder if this is some limitation of mount or if this is OpenVZ related.

```
hardwarenode:~# df -h
```

```
/dev/mapper/dataone 250M 97M 140M 41% /data1
```

```
/dev/mapper/datatwo 250M 97M 140M 41% /data2
```

```
hardwarenode:~# mount
```

```
/dev/mapper/dataone on /data1 type ext3 (rw)
```

```
/dev/mapper/datatwo on /data2 type ext3 (rw)
```

```
/data1/somename on /vz/root/105/site/dataone type none (rw,bind)
```

```
/data2/somename on /vz/root/105/site/datatwo type none (rw,bind)
```

```
hardwarenode:~# vzctl enter 105
```

```
entered into VE 105
```

```
somevps:~# mount
```

```
simfs on / type simfs (rw)
```

```
proc on /proc type proc (rw)
```

```
sysfs on /sys type sysfs (rw)
```

```
devpts on /dev/pts type devpts (rw)
```

```
tmpfs on /dev/shm type tmpfs (rw)
```

```
ext3 on /site/dataone type ext3 (rw,data=ordered)
```

```
ext3 on /site/datatwo type ext3 (rw,data=ordered)
```

```
ext3 on /site/dataone/ftpd/ftp/site/1 type ext3 (rw,data=ordered)
```

```
ext3 on /site/dataone/ftpd/ftp/site/2 type ext3 (rw,data=ordered)
```

```
ext3 on /site/dataone/ftpd/ftp/site/3 type ext3 (rw,data=ordered)
```

```
ext3 on /site/dataone/ftpd/ftp/site/4 type ext3 (rw,data=ordered)
```

```
ext3 on /site/dataone/ftpd/ftp/site/5 type ext3 (rw,data=ordered)
```

```
ext3 on /site/dataone/ftpd/ftp/site/6 type ext3 (rw,data=ordered)
```

```
somevps:~# df -h
```

```
Filesystem      Size  Used Avail Use% Mounted on
```

```
simfs           1.9G 356M 1.6G 19% /
```

```
tmpfs           1012M  0 1012M  0% /dev/shm
```

```
ext3            459G  29G 407G  7% /site/dataone
```

```
ext3            459G  8.2G 428G  2% /site/datatwo
```

```
ext3            459G  8.2G 428G  2% /site/dataone/ftpd/ftp/site/1
```

```
ext3            459G  8.2G 428G  2% /site/dataone/ftpd/ftp/site/2
```

```
ext3          459G  8.2G  428G  2% /site/dataone/ftpd/ftp/site/3
ext3          459G  8.2G  428G  2% /site/dataone/ftpd/ftp/site/4
ext3          459G  8.2G  428G  2% /site/dataone/ftpd/ftp/site/5
ext3          459G  8.2G  428G  2% /site/dataone/ftpd/ftp/site/6
```

Thanks in advance,

Subject: Re: Mounts--bind, issue or limitation?

Posted by [Vasily Tarasov](#) on Mon, 07 May 2007 07:07:22 GMT

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

Hello,

I want to clear up some ambiguities:

1. You wrote, that you had two 500GB disk drives, but df on HN outputs:

```
hardwarenode:~# df -h
/dev/mapper/dataone 250M 97M 140M 41% /data1
/dev/mapper/datatwo 250M 97M 140M 41% /data2
```

So, what are dataone and datatwo? Partions of the same drive or different drives? Or do I miss something?

2. Why in VE df outputs other size (other then 250M) of filesystems on /dataone /datatwo partions:

```
somevps:/# df -h
...
ext3 459G 29G 407G 7% /site/dataone
ext3 459G 8.2G 428G 2% /site/datatwo
...
```

3. In VE mount command outputs:

```
....
ext3 on /site/dataone type ext3 (rw,data=ordered)
ext3 on /site/datatwo type ext3 (rw,data=ordered)
ext3 on /site/dataone/ftpd/ftp/site/1 type ext3 (rw,data=ordered)
ext3 on /site/dataone/ftpd/ftp/site/2 type ext3 (rw,data=ordered)
ext3 on /site/dataone/ftpd/ftp/site/3 type ext3 (rw,data=ordered)
ext3 on /site/dataone/ftpd/ftp/site/4 type ext3 (rw,data=ordered)
ext3 on /site/dataone/ftpd/ftp/site/5 type ext3 (rw,data=ordered)
ext3 on /site/dataone/ftpd/ftp/site/6 type ext3 (rw,data=ordered)
```

Bind mount points /site/dataone/ftpd/ftp/site/1, /site/dataone/ftpd/ftp/site/2, ... are mount points from what locations? I want to know that in order to probably reproduce the situation here on our local node.

4. The last but probably the most important. You wrote, that "This resulted in me not being able to

use the disks at all on the hardware node". What does that mean? How exactly you can't use your drives? Can you post here exact commands that fail to work?

Thank you,
Vasily

Subject: Re: Mounts--bind, issue or limitation?
Posted by [3molo](#) on Mon, 07 May 2007 11:28:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

On the hardware node the disks was usable and I could see their capacity with df. Now df claims them to be 250MB filesystems which is NOT the case. ls shows me nothing, although these are the very disks that are mount binded into the vps. How and why, I dont know. It's a very weird situation. There can be nothing I'm reading wrong, since I have one 18gb scsi + 9gb scsi + 2x 500 gb ide disks. Also they worked on the HN when I first mounted them, then I discovered the missinformation you can see below, and of course the many mounts. The mount binds the owner of the vps did was from (inside his VE) /site/dataone/{1,2,3} + /site/datatwo/{4,5,6} into /site/dataone/ftpd/ftp/site/.

Clearly not true information here:

```
hardwarenode:~# df -h
/dev/mapper/dataone 250M  97M 140M 41% /data1
/dev/mapper/datatwo 250M  97M 140M 41% /data2
```

```
hardwarenode:~# mount
/dev/mapper/dataone on /data1 type ext3 (rw)
/dev/mapper/datatwo on /data2 type ext3 (rw)
/data1/midimiliz on /vz/root/105/site/dataone type none (rw,bind)
/data2/midimiliz on /vz/root/105/site/datatwo type none (rw,bind)
```

```
hardwarenode:~# ls -la /data{1,2}
/data1:
total 2
drwxr-xr-x  2 root root 1024 2007-04-27 14:16 .
drwxr-xr-x 25 root root 1024 2007-04-27 14:40 ..
```

```
/data2:
total 2
drwxr-xr-x  2 root root 1024 2007-04-18 13:04 .
drwxr-xr-x 25 root root 1024 2007-04-27 14:40 ..
```

```
hardwarenode:~# du -sh /vz/root/105/site/data{one,two}
37G  /vz/root/105/site/dataone
8.0G  /vz/root/105/site/datatwo
```

```
somevps:~# df -h
ext3          459G  29G 407G  7% /site/dataone
ext3          459G  8.2G 428G  2% /site/datatwo
```

So only way for me to read or seek the disks on th HN is by using the mount point /vz/root/105/site/data{one,two}. Im pretty sure that without the many mount binds inside the vps, there would be no problem. But I would really like to know why and what to do before umounting their structure!

Hope I made the problem more clear

Subject: Re: Mounts--bind, issue or limitation?
Posted by [Vasily Tarasov](#) on Mon, 07 May 2007 12:20:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

Now I get the problem completely. Really weird situation!
Well, first of all I'll try to reproduce it here on my node. I'll inform you soon about results.

Thank you,
Vasily.

Subject: Re: Mounts--bind, issue or limitation?
Posted by [Vasily Tarasov](#) on Mon, 07 May 2007 12:26:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

And what kernel version you're using?

Thanks.

Subject: Re: Mounts--bind, issue or limitation?
Posted by [3molo](#) on Mon, 07 May 2007 12:46:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

Linux goebbels 2.6.18-028stab027 #1 SMP Fri Apr 27 14:10:13 CEST 2007 i686 GNU/Linux
mount: mount-2.12r

You give me some hope, Vasily

Subject: Re: Mounts--bind, issue or limitation?
Posted by [Vasily Tarasov](#) on Tue, 08 May 2007 10:16:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

I conducted some experiments on my node and I also had some strange symptoms with bind mounts in VE. Can you, please, fill the bug in bugzilla with detailed description?

Thank you,
Vasily.

Subject: Re: Mounts--bind, issue or limitation?
Posted by [3molo](#) on Tue, 08 May 2007 11:33:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

So this is OpenVZ related?

Subject: Re: Mounts--bind, issue or limitation?
Posted by [Vasily Tarasov](#) on Wed, 09 May 2007 06:40:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

I can't say it exactly now, it can be a mainstream problem. But even in this case - fill the bug in OpenVZ bugzilla. If later it'll turn out to be a mainstream problem - we'll port the fix.

Thank you,
Vasily.
