

My thanks to vass for pointing me in the right direction. I will give my solution and then discuss documentation.

The basic problem was that the hardware node (node0) treats SATA drives as SCSI disks. The same problem no doubt exists with real SCSIs.

After formatting the sda drive and creating filesystems from node0, my attempts mount sda1, etc. from ANY of the nodes failed, with the message, "device busy". By this time I had already given access by vzctl set 150 --devices 8:0:rw 8:1:rw, etc. to the sda drive. I found the answer in /proc/partitions on node0, which showed that the drive partitions had been mapped to /dev/dm-0, /dev/dm-1, etc. Apparently device-mapper runs when node1 boots.

Now I have no interest in handling the SATA drive as a logical volume, so my first step was to get rid of the mapping. I discovered dmsetup and, after making sure sda was the only mapped device, I ran it:

```
/sbin/dmsetup remove_all
```

Then I was able to mount sda1 on node1. Following the advice from sass, I used vzctl set 151 --devnodes sda rw --save.

Going to node 150 I discovered I still had a problem; the mount command gave an error:

```
mount: wrong fs type, bad option, bad superblock on /dev/sda1,  
       or too many mounted file systems
```

I believe this error happened because I originally ran mkfs from node0. I could probably have corrected it on node 150 by running mkfs on the partition there. Instead, I went back to node0 and created a mount path there, and then did a mount --bind to a directory in node 150's private area. Once this was done, I could mount sda partitions in node 150.

Note that some virtual nodes do not have many devices in /dev, and they may have to be created using mknod.

Bottom line: there are two solutions to mounting an unshared partition on a VPS:

1. Get authorization from node0 to partition the drive directly from the VPS and create its filesystems there.
2. Mount the partition on node0 and use mount --bind to link it to a mount directory on the VPS.

Briefly about documentation.

Usually I try to search for technical info before asking help from a forum. In this case I missed the fact that --devnodes is an option under vzctl. The User's Guide does not discuss this subject at all, though it does mention -devices briefly. Also none of the FAQs deal with non-quota disk management in VPSs.

Sys admins who need to create specialized VPSs have to have more information about setting up VPSs in non-standard ways. This situation can be expected in new technology. Hopefully, after I gain more experience with openvz, I can help out by writing some new FAQs or HOWTOs.

Thanks again,
rollinw
