
Subject: NFS Performance

Posted by [dlight](#) on Thu, 22 Sep 2011 18:14:15 GMT

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

First and foremost, this is my first foray into OpenVZ and I have to say I'm loving it so far. (Coming from Solaris Containers) That being said, I'm not that versed in the nuances yet.

What I have is what I believe to be a very straightforward config from a networking and resource perspective. I didn't set any VZ limits and haven't hit a single "failcnt" in user_beancounters.

I have a veth device assigned to the VZ instance and that virtual device interface is joined to a bridge on the HW node. (1Gb/Full HW link speed) Networking works just fine...

However, when conducting very basic Bandwidth tests over NFS I am getting some horrible numbers while using `cp`, but not with other utilities (dd, rsync, etc). And this only occurs when running a `cp` in a VZ instance.

Ex Data:

VZ:

NFS Mount Options:

```
rw,nfsvers=3,rsize=32768,wsiz=32768,acregmin=0,acregmax=0,acdirmin=0,acdirmax=0,hard,proto=tcp,timeo=600,retrans=2,addr=na1
```

```
[root@VZ /]# time cp /s55/test/1GB ./
```

```
real 1m12.103s
user 0m0.269s
sys 0m23.979s
```

```
[root@VZ /]# time rsync -vpa /s55/test/1GB ./
```

```
building file list ... done
1GB
```

```
sent 1024125088 bytes received 42 bytes 81930010.40 bytes/sec
total size is 1024000000 speedup is 1.00
```

```
real 0m11.772s
user 0m12.355s
sys 0m7.454s
```

HW Node:

NFS Mount Options:

```
rw,nfsvers=3,rsize=32768,wsiz=32768,acregmin=0,acregmax=0,acdirmin=0,acdirmax=0,hard,proto=tcp,timeo=600,retrans=2,addr =na1
```

```
[root@HWNODE /]# time cp /s55/test/1GB ./
```

```
real 0m9.543s  
user 0m0.006s  
sys 0m0.941s
```

```
[root@HWNODE /]# time rsync -vpa /s55/test/1GB ./  
building file list ... done  
1GB
```

```
sent 100884596 bytes received 42 bytes 67256425.33 bytes/sec  
total size is 100872192 speedup is 1.00
```

```
real 0m1.147s  
user 0m1.179s  
sys 0m0.657s
```

I get similar results to rsync while using dd....

So what am I missing about "how" `cp` moves data, or do I need to configure a different scheduler?

All in all, I wouldn't care to much about `cp` performance, but I am seeing very similar behavior in Oracle to what I am seeing with cp. (I can't obviously switch out Oracle for rsync no matter how nice it sounds...)

Anyone have any ideas for me?
