
Subject: Re: VZ+DRBD+HA and load balancing
Posted by [golly](#) on Wed, 27 Aug 2008 14:36:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

This is really quick and rough to give you an idea of what I did. I'll try to follow it up once I have all my notes arranged. I really *JUST* got done setting this up (as in I got it fully running *yesterday*...).

In laying out my directories, I wanted to set the nodes up so that I could recreate the nodes as quickly as possible in the event of a hardware failure. Therefore, I split files into those on a default install of Ubuntu server and those that I modified or that needed to be replicated between nodes.

So, once I had an install of Ubuntu Hardy server running, I set about using the /srv directory as the place where all of my "replication files" went. I mounted a 1G block device at /srv and will be setting up file replication with rsync for files *on that block device only* soon. I didn't want to use DRBD for the /srv block device because it has drbd.conf and heartbeat config files.

Anyway, my current /srv filesystem looks like this (not including files copied over from /etc or /var of the default install of drbd/heartbeat/opensvz):

```
/srv
|-/heartbeat
| |-authkeys
| |-ha.cf
|
|_-/opensvz
| |-/conf
| | |-1008.conf -> ../ve/1008/opensvz-ve.conf
| | |-1009.conf -> ../ve/1009/opensvz-ve.conf
| |
| |-/cron
| |-/dists
| |-/dump
| |-/lock
| |-/names
| |
| |-/templates
| | |-/cache
| | |-ubuntu-8.04-i386-minimal.tar.gz
| |
| |-/ve
| | |-/1008
| | |-/1009
|
|-drbd.conf
|-sensors3.conf
```

NOTES:

1. /srv/opensvz/ve/{1008,1009,...} are the mount points for the individual virtual machine block devices.
2. /etc/vz is the only *directory* symlink (symlinked to /srv/opensvz), the rest are file symlinks (i.e. drbd.conf, heartbeat's ha.cf and authkeys, etc.)

The virtual machine block devices have this directory layout:

```
/dump
/private
/root
opensvz-ve.conf
```

NOTE: I will eventually use the VE's /dump to handle live-migration, but its not there yet....

To create a new virtual machine using this layout, follow this outline:

1. create a new DRBD device on two nodes and start sync'ing them.
2. make one side primary and create filesystems.
3. create VE mount point (i.e. /srv/opensvz/ve/\$VEID).
4. mount primary DRBD device on VE mount point.
5. vzctl create....
You can set up the vm here or wait until its running under heartbeat##
6. copy /etc/vz/conf/\$VEID.conf to VE mount point and symlink it back to /etc/vz/conf/\$VEID.conf.
7. unmount primary DRBD device.
8. wait for DRBD to finish sync'ing and then do "drbdadm down" on both nodes.
9. search 'n replace the constraints and resources XML with your info.
10. sudo cibadmin -C -x vm-constraints.xml -o constraints
11. sudo cibadmin -C -x vm-resources.xml -o resources
12. sudo crm_resource -r ms-##DRBDDEVICE## --meta -p target-role -v "#default"
13. sudo crm_resource -r gp-##VENAME## --meta -p target-role -v "#default"

I've attached all of the config files and heartbeat "template" XML files that I think are central to my setup. If I missed anything, let me know....

Finally, about heartbeat V1 - sorry, I wanted V2 as I don't know if the ManageVE ocf script works with V1. And, while we are talking about the ManageVE script - do yourself a favor and change the #!/bin/sh to #!/bin/bash at the top of /usr/lib/ocf/resource.d/heartbeat/ManageVE.

I can have more detailed instructions for node setup and vm setup on Ubuntu Hardy if there is enough interest.

Enjoy!

File Attachments

- 1) [drbd.conf](#), downloaded 426 times
- 2) [vz.conf](#), downloaded 397 times
- 3) [opensvz-ve.conf](#), downloaded 438 times

- 4) [vm-constraints.xml](#), downloaded 401 times
 - 5) [vm-resources.xml](#), downloaded 391 times
-