

Building Elastix-2.3 High Availability Clusters with DRBD and Heartbeat

Disclaimer

Use the information in this document at your own risk. Use of the concepts, examples, and/or other content of this document are entirely at your own risk. All copyrights are owned by their owners, unless specifically noted otherwise. Use of a term in this document should not be regarded as affecting the validity of any trademark or service mark. You are strongly recommended to take a backup of your system before major installation and backups at regular intervals.

Credits

Special thanks to Telesoft Integrando Technologies whose earlier documentation on doing DRBD in Elastix was a great reference for this work;

<http://asterisk.aplitel.info/files/asteriskcluster.pdf>

Another special thanks to **Amjad Jabali** for the Elastix 1.6.0 Clustering tutorial that this document is based on.

<http://ftp.heanet.ie/disk1/sourceforge/e/el/elastix-easy/Documents/Elastix-Clustering.pdf>

A final special thanks to the Escuela Politécnica del Ejercito (Sangolquí – Ecuador) for the help provided by their staff and faculty members in the Elastix Project Overview.

If by any way you can contribute to polish this work feel free to write to my e-mail and send me your documentation. Remember Open Source is based on community effort, so any recommendation to improve it is appreciated and valued.

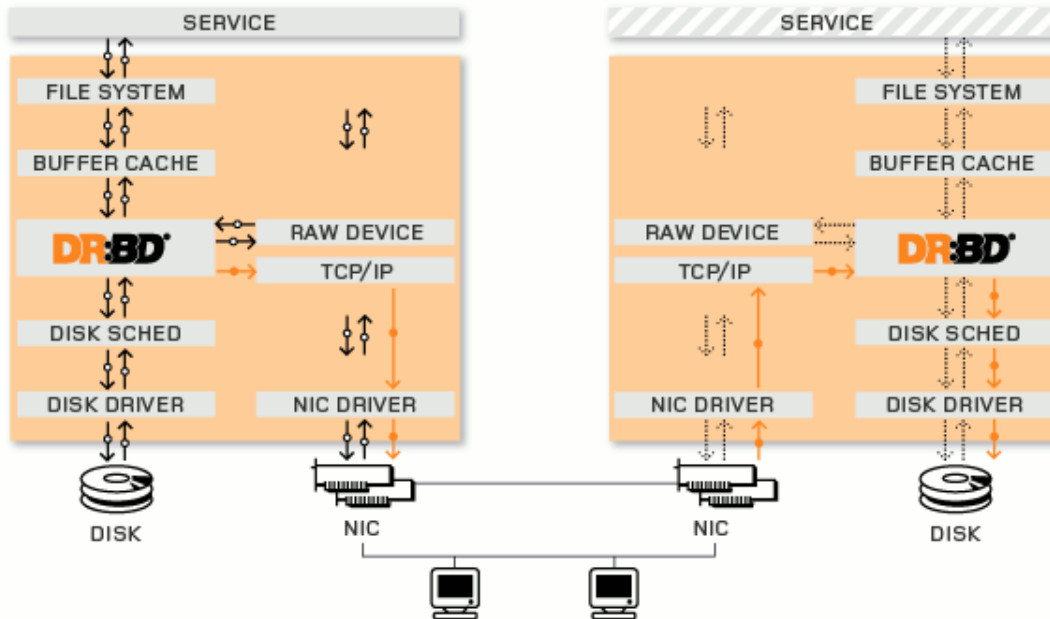


INDEX

Operational Overview.....	3
What Is DRBD.....	3
What Heartbeat does.....	4
Equipment Overview.....	4
DRBD Install and Configuration.....	5
Heartbeat Configuration.....	10
Credits.....	12
References.....	12

Operational Overview

What is DRBD?



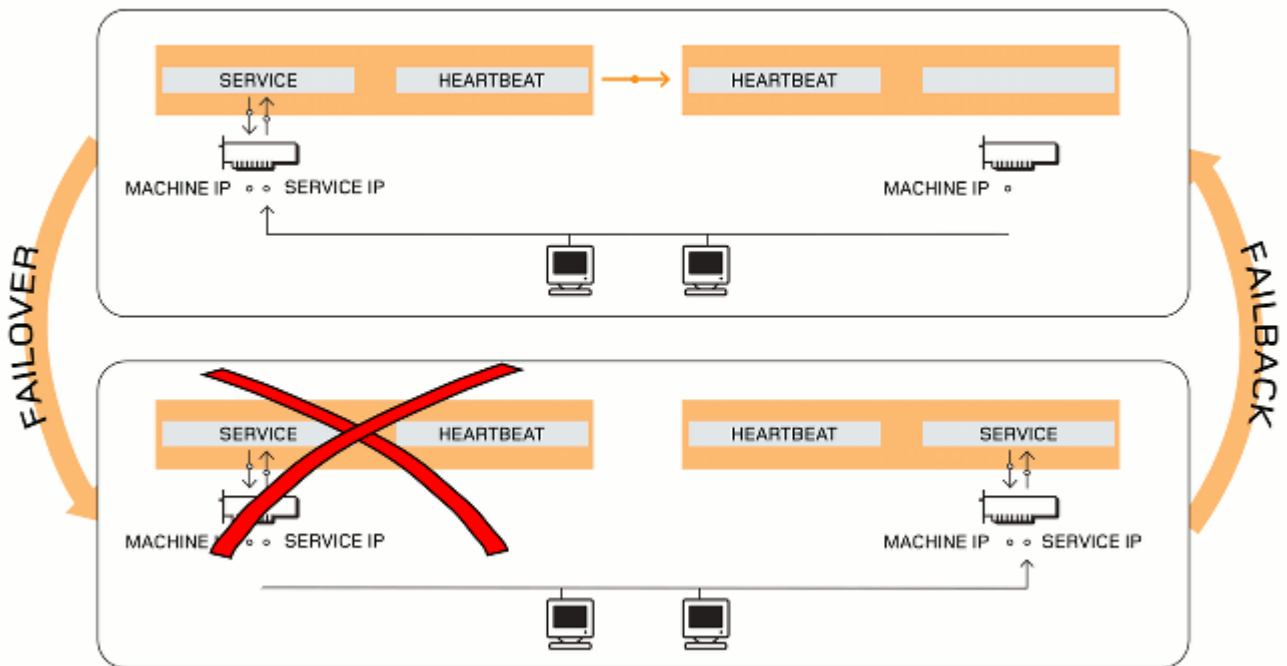
DRBD[®] refers to block devices designed as a building block to form high availability (HA) clusters. This is done by mirroring a whole block device via an assigned network. DRBD can be understood as network based raid-1.

In the illustration above, the two orange boxes represent two servers that form an HA cluster. The boxes contain the usual components of a Linux[™] kernel: file system, buffer cache, disk scheduler, disk drivers, TCP/IP stack and network interface card (NIC) driver. The black arrows illustrate the flow of data between these components.

The orange arrows show the flow of data, as DRBD mirrors the data of a highly available service from the active node of the HA cluster to the standby node of the HA cluster.

In our implementation we will be creating a DRBD synchronized partition on `/dev/sda3` called "replica". This partition will contain only those directories and files we want synchronized between our primary and secondary server. Namely, the important Asterisk and Elastix related directories and files.

What Heartbeat does




The upper part of this picture shows a cluster where the left node is currently active, i.e., the service's IP address that the client machines are talking to is currently on the left node.

The service, including its IP address, can be migrated to the other node at any time, either due to a failure of the active node or as an administrative action. The lower part of the illustration shows a degraded cluster. In HA speak the migration of a service is called *failover*, the reverse process is called *failback* and when the migration is triggered by an administrator it is called *switchover*.

In our implementation we will utilize Heartbeat to monitor the state of two servers and during a failover mount our synchronized partition on the secondary server and start up the following resources/applications; asterisk, mysql and http. During failover our floating IP address will move from the primary to the secondary server. This IP address should be used to register SIP and other VoIP endpoints.

Equipment Overview

This installation scenario assumes two servers, each with three Ethernet interfaces and a single SATA hard drive. You may have a different type of hard drive (IDE, SCSI, etc) and therefore some of these steps may need to be modified to better reflect your environment.



DRBD Install and Configuration

The following steps are to be performed on both primary and secondary servers:

1. Boot Elastix-2.3 Install CD
2. From boot menu, type "advanced" and enter.
3. During install routine choose to manually partition hard drive. The following is based on a 500.0 GB SATA with 4GB RAM.
 - Create root (/), ext3 partition with 50000MB (sda1) (give 10% of your hard drive for future growth).
 - Create swap partition with 8072MB (sda2) (choose your swap with double the RAM size).
4. When prompted to enter ip addresses, enter the ones needed for your implementation and change the localhost names to voipserver.drbd and voipbackup.drbd. The remainder of the install routine is standard.
5. After installation and booting perform upgrade

```
yum -y update
```
6. Create partition that will contain the replicated data
 - fdisk /dev/sda
 - Add a new partition (n)

- Primary (p)
- Partition number (3)
- Press enter until returned to fdisk command prompt
- **NOTE:** if your servers have two different sized hard drives it is imperative that the third partition is identical in size or they will never synchronize over DRBD. Do this by accepting the default first cylinder and then specifying the Last cylinder with the +sizeM option. Ex. +6048M. Make these same specifications on both servers.
- Press “t” to change the partition system ID
- Press “3” to choose partition number
- Choose HEX 83 for type
- Press “w” to save changes

RESTART SERVER

7. Format newly made partition

```
mke2fs -j /dev/sda3
```

8. Now we delete the file system from the disk we just created

```
dd if=/dev/zero bs=1M count=1 of=/dev/sda3; sync
```

9. Install DRBD, Heartbeat and dependencies with yum.

```
yum install heartbeat drbd83 kmod-drbd83
```

Note: If by any chance you experience problems with drbd83, use drbd82 version (64 bit versions).

10. To ensure proper host name to IP resolution it is recommended that you manually update the /etc/hosts file to reflect proper host-to-IP mapping.

```
192.168.1.243 voipserver.drbd
192.168.1.242 voipbackup.drbd
127.0.0.1      localhost voipbackup.drbd localhost.localdomain
::1          localhost6.localdomain6 localhost6
```

11. Edit /etc/drbd.conf on Server1.drbd. Modify this sample to meet your particular needs.

```
global { usage-count no; }
resource r0 {
protocol C;
startup { wfc-timeout 10; degr-wfc-timeout 30; } #change timers
to your need
disk { on-io-error detach; } # or panic, ...
net {
```

```

        after-sb-0pri discard-younger-primary; #change options
to your need
        after-sb-1pri discard-secondary;
        after-sb-2pri call-pri-lost-after-sb;
        cram-hmac-alg "sha1";
        shared-secret "Cent0Sru!3z";
    }

    syncer { rate 100M; }
    on voipserver.drbd {
        device /dev/drbd0;
        disk /dev/sda3;
        address 10.1.1.1:7788;
        meta-disk internal;
    }
    on voipbackup.drbd {
        device /dev/drbd0;
        disk /dev/sda3;
        address 10.1.1.2:7788;
        meta-disk internal;
    }
}

```

Note: The following lines are used to help the servers resolve split brain recovery. Split brain is when two servers are in primary mode and need to know how to resolve who should assume primary/secondary role (discarding or accepting changes made in primaries).

```

    after-sb-0pri discard-younger-primary;
    after-sb-1pri discard-secondary;
    after-sb-2pri call-pri-lost-after-sb;

```

Reference: <http://www.drbd.org/users-guide/s-configure-split-brain-behavior.html>

12. Replicate this config file (/etc/drbd.conf) to the second server

```
scp /etc/drbd.conf root@voipbackup.drbd:/etc/
```

13. Initialize the meta-data area on disk before starting drbd (! on both server!)

```
drbdadm create-md r0
```

14. Start drbd on both nodes (service drbd start)

```
service drbd start
```

15. Verify that both server are secondary

```
cat /proc/drbd
```

16. As you can see, both nodes are secondary, which is normal. we need to decide which node will act as a primary now (voipserver.drbd) : that will initiate the first 'full sync' between the two nodes:

```
drbdadm -- --overwrite-data-of-peer primary r0
```

17. Launch the command and wait until it's finish synchronizing

```
watch -n 1 cat /proc/drbd
```

18. We can now format /dev/drbd0 and mount it on voipserver.drbd:

```
mkfs.ext3 /dev/drbd0  
mkdir /replica  
mount /dev/drbd0 /replica
```

19. We can determine the role of a server by executing the following;

```
drbdadm role r0
```

The primary server should return;

```
Primary/Secondary
```

20. Now we will copy all of the directories we want synchronized between the two servers to our new partition, remove the original directories and then create symbolic links to replace them on voipserver.drbd

```
cd /replica  
tar -zcvf etc-asterisk.tgz /etc/asterisk  
tar -zxvf etc-asterisk.tgz  
tar -zcvf var-lib-asterisk.tgz /var/lib/asterisk  
tar -zxvf var-lib-asterisk.tgz  
tar -zcvf usr-lib-asterisk.tgz /usr/lib/asterisk/  
tar -zcvf var-www.tgz /var/www/  
tar -zxvf usr-lib-asterisk.tgz  
tar -zcvf var-spool-asterisk.tgz /var/spool/asterisk/  
tar -zxvf var-spool-asterisk.tgz  
tar -zcvf var-lib-mysql.tgz /var/lib/mysql/  
tar -zxvf var-lib-mysql.tgz  
tar -zcvf var-log-asterisk.tgz /var/log/asterisk/  
tar -zxvf var-log-asterisk.tgz  
tar -zxvf var-www.tgz  
rm -rf /etc/asterisk  
rm -rf /var/lib/asterisk  
rm -rf /usr/lib/asterisk/  
rm -rf /var/spool/asterisk
```



```

rm -rf /var/lib/mysql/
rm -rf /var/log/asterisk/
ln -s /replica/etc/asterisk/ /etc/asterisk
ln -s /replica/var/lib/asterisk/ /var/lib/asterisk
ln -s /replica/usr/lib/asterisk/ /usr/lib/asterisk
ln -s /replica/var/spool/asterisk/ /var/spool/asterisk
ln -s /replica/var/lib/mysql/ /var/lib/mysql
ln -s /replica/var/log/asterisk/ /var/log/asterisk
ln -s /replica/var/www /var/www

```

Note: This worked and replicated perfectly. If you think a new directory should be added for this new version or something is missing, feel free to write to my e-mail with the suggestions. This tutorial is updated from the 1.6 HA Cluster tutorial from the Elastix Forums.

21. Stop mysqld, asterisk and httpd services on voipserver.drbd

```
service mysqld restart
```

```

service mysqld stop
service asterisk stop
service httpd stop

```

22. Verify services are down and proceed to switch manually to the second server:

```
[root@voipserver.drbd /]# umount /replica ; drbdadm secondary r0
```

```

[root@voipbackup.drbd /]# mkdir /replica ; drbdadm primary r0 ; mount /dev/drbd0 /replica
[root@voipbackup.drbd /]# ls /replica/

```

Note: This is used to check if you are replicating information on both servers. You should see all data replicated in the secondary server just like data in the primary.

23. Verify voipserver.drbd status (Primary/Secondary)

```
drbdadm role r0
```

24. Execute 'df -h' on the primary to confirm that our /dev/drbd0 partition is mounted and in use.

```

[root@voipserver ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda1        48G  3.8G  42G   9% /
tmpfs            1.9G   0  1.9G   0% /dev/shm
/dev/drbd0       403G  640M  382G   1% /replica

```

Note: Executing this same command in voipbackup.drbd while in secondary mode should not display the /dev/drbd0 partition unless it's assuming primary mode.

25. Now we will remove and link on voipbackup.drbd

```

rm -rf /etc/asterisk
rm -rf /var/lib/asterisk
rm -rf /usr/lib/asterisk/

```

```
rm -rf /var/spool/asterisk
rm -rf /var/lib/mysql/
rm -rf /var/log/asterisk/
ln -s /replica/etc/asterisk/ /etc/asterisk
ln -s /replica/var/lib/asterisk/ /var/lib/asterisk
ln -s /replica/usr/lib/asterisk/ /usr/lib/asterisk
ln -s /replica/var/spool/asterisk/ /var/spool/asterisk
ln -s /replica/var/lib/mysql/ /var/lib/mysql
ln -s /replica/var/log/asterisk/ /var/log/asterisk
ln -s /replica/var/www /var/www
```

26. Stop mysqld, asterisk and httpd services on voipbackup.drbd

```
service mysqld restart
```

```
service mysqld stop
service asterisk stop
service httpd stop
```

27. Now switch back to the first server :

```
[root@voipbackup.drbd /]# umount /replica/ ; drbdadm secondary r0
[root@voipserver.drbd /]# drbdadm primary r0 ; mount /dev/drbd0 /replica
```

28. Drbd is working ... let's be sure that it will always be started:

```
chkconfig drbd83 on
```

Heartbeat Configuration

29. Remember to stop any boot up services on both servers that should be controlled by heartbeat. These services will be controlled by heartbeat on the server that is in control.

```
chkconfig asterisk off
chkconfig mysqld off
chkconfig httpd off
service mysqld stop
service asterisk stop
service httpd stop
```

30. Let's configure a simple /etc/ha.d/ha.cf file on voipserver.drbd :

```
debugfile /var/log/ha-debug
logfile /var/log/ha-log
```

```
logfacility local0
keepalive 2
deadtime 30
warntime 10
initdead 120
udpport 694
bcast eth1
auto_failback on
node voipserver.drbd
node voipbackup.drbd
```

31. Create also the /etc/ha.d/authkeys on voipserver.drbd:

```
auth 1
1 sha1 MySecret
```

32. Change permissions on the /etc/ha.d/authkeys file on voipserver.drbd:

```
chmod 600 /etc/ha.d/authkeys
```

33. Edit /etc/ha.d/haresources on voipserver.drbd: (It is one line!!!!!! Formating is important)

```
voipserver.drbd drbddisk::r0 Filesystem::/dev/drbd0::/replica::ext3 IPaddr::192.168.1.245/24/eth0/192.168.1.255 mysqlid asterisk httpd
```

34. Start the heartbeat service on voipserver.drbd :

```
service heartbeat start
```

35. Replicate now the ha.cf, authkeys and haresources to voipbackup.drbd and start heartbeat

```
[root@voipserver.drbd ha.d]# scp /etc/ha.d/ha.cf /etc/ha.d/authkeys /etc/ha.d/haresources
root@voipbackup.drbd:/etc/ha.d/
```

```
[root@svoipbackup.drbd ha.d]# service heartbeat start
```

36. Configure heartbeat to initialize at boot on both server

```
chkconfig --add heartbeat
```

37. Verify voipserver.drbd status (Primary/Secondary)

```
drbdadm role r0
```

38. Execute 'df -h' on the primary to confirm that our /dev/drbd0 partition is mounted and in use.

```
[root@voipserver ~]# df -h
Filesystem      Size Used Avail Use% Mounted on
/dev/sda1       48G  3.8G  42G   9% /
tmpfs           1.9G  0  1.9G   0% /dev/shm
/dev/drbd0      403G  640M  382G   1% /replica
```

39. Test your work by creating a SIP extension or anything inside Elastix Web Interface, then shut down your primary server while making a continuous ping to 192.168.1.245 (floating IP address) verifying it doesn't lose connectivity. Make another change in the secondary server, turn your primary back on, and all changes should be kept intact.

Special Note: Any changes made to asterisk files should be done via web Interface ONLY. Do not attempt to upgrade Elastix version once finished the cluster or else it will write its own files again discarding links to the /replica directory.

Troubleshooting:

```
tcpdump -i eth0:0 -s 1500 -w captura.pcap          #capture traffic
mv captura.pcap /var/www/html                     #move file to web for download
```

Credits

MaxiDistribuciones Cía. Ltda.
www.grupomaxi.com.ec

Escuela Politécnica del Ejército
www.espe.edu.ec

Telesoft Integrando Technologies

Redfone

References

<http://wiki.centos.org/HowTos/Ha-Drbd>
http://support.red-fone.com/downloads/elastix/Elastix_HA_Cluster.pdf
<http://danielaliaman.com/blog/files/phonecube/cluster/AsteriskCluster.pdf>
<http://www.drbd.org/users-guide/s-configure-split-brain-behavior.html>

Author: Daniel Guevara
E-mail: danielguevara14@hotmail.com

Last Updated: 2-May-12