By Brian Hellman
Published: 2009-02-27 15:23

DRBD 8.3 Third Node Replication With Debian EtchInstallation and Set Up Guide for DRBD 8.3 + Debian EtchThe Third Node Setup

by Brian Hellman

The recent release of DRBD 8.3 now includes ***The Third Node*** feature as a freely available component. This document will cover the basics of setting up a third node on a standard Debian Etch installation. At the end of this tutorial you will have a DRBD device that can be utilized as a SAN, an iSCSI target, a file server, or a database server.

Note: LINBIT support customers can skip Section 1 and utilize the package repositories.

LINBIT has hosted third node solutions available, please contact them at sales_us at linbit.com for more information.

Preface:

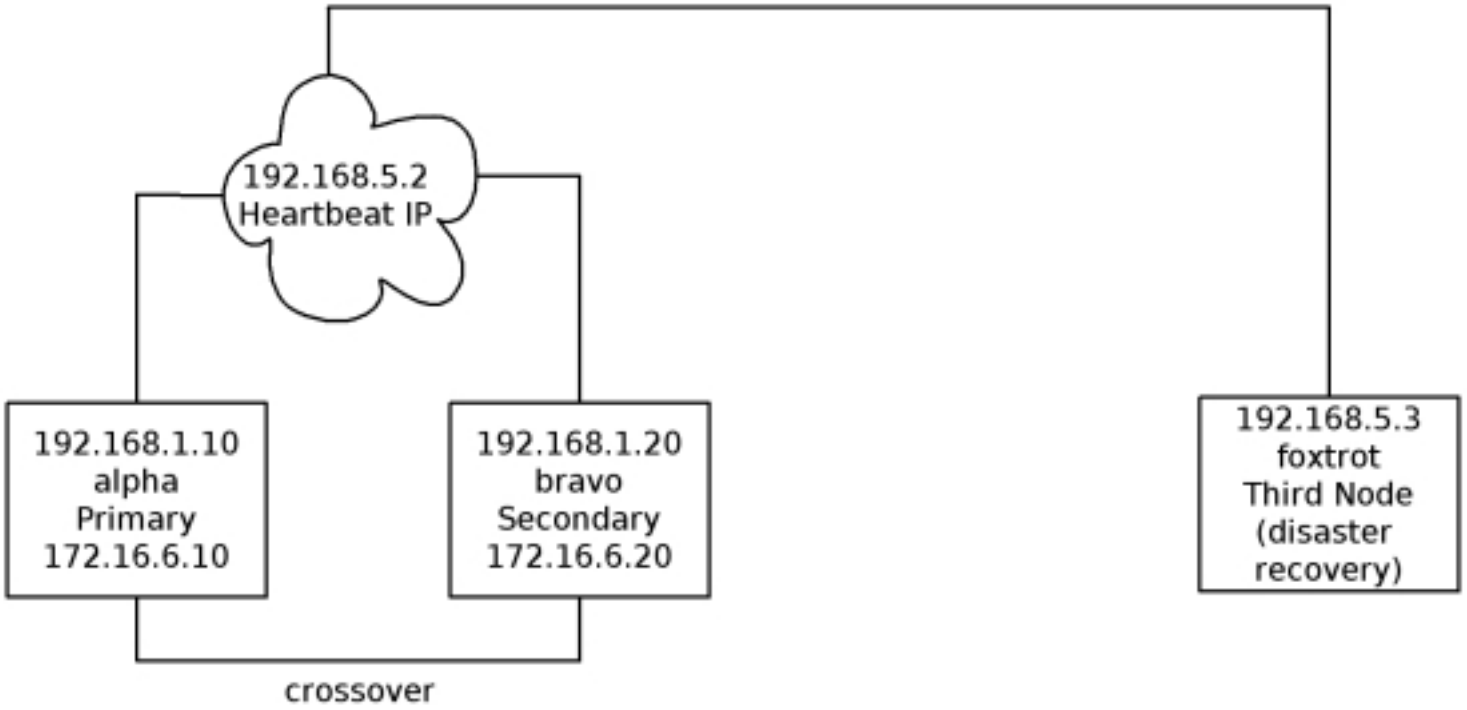

The setup is as follows:

- Three servers: *alpha*, *bravo*, *foxtrot*
- *alpha* and *bravo* are the primary and secondary local nodes
- *foxtrot* is the third node which is on a remote network
- Both *alpha* and *bravo* have interfaces on the *192.168.1.x* network (*eth0*) for external connectivity.
- A crossover link exists on *alpha* and *bravo* (*eth1*) for replication using *172.16.6.10* and *.20*
- Heartbeat provides a virtual IP of *192.168.5.2* to communicate with the disaster recovery node located in a geographically diverse location

Section 1: Installing The Source

These steps need to be done on each of the 3 nodes.

Prerequisites:

- make
- gcc
- glibc development libraries
- flex scanner generator
- headers for the current kernel

Enter the following at the command line as a privileged user to satisfy these dependencies:

```
apt-get  install make gcc libc6 flex linux-headers-`uname -r` libc6-dev linux-kernel-headers
```

Once the dependencies are installed, download DRBD. The latest version can always be obtained at ***http://oss.linbit.com/drbd/***. Currently, it is 8.3.

```
cd /usr/src/

wget http://oss.linbit.com/drbd/8.3/drbd-8.3.0.tar.gz
```

After the download is complete:

- Uncompress DRBD
- Enter the source directory
- Compile the source
- Install DRBD

```
tar -xzvf drbd-8.3.0.tar.gz

cd /usr/src/drbd-8.3.0/
```

```
make clean all

make install
```

Now load and verify the module:

```
modprobe drbd

cat /proc/drbd
```

```
  version: 8.3.0 (api:88/proto:86-89)
GIT-hash: 9ba8b93e24d842f0dd3fb1f9b90e8348ddb95829 build by root@alpha, 2009-02-05 10:36:11
```

Once this has been completed on each of the three nodes, continue to next section.

Section 2: Heartbeat Configuration

Setting up a third node entails stacking DRBD on top of DRBD. A virtual IP is needed for the third node to connect to, for this we will set up a simple Heartbeat v1 configuration. This section will only be done on `alpha` and `bravo`.

Install Heartbeat:

```
apt-get install heartbeat
```

Edit the ***authkeys*** file:

```
vi /etc/ha.d/authkeys
```

```
auth 1
1 sha1 yoursupersecretpasswordhere
```

Once the file has been created, change the permissions on the file. Heartbeat will not start if this step is not followed.

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

Copy the ***authkeys*** file to *bravo*:

```
scp /etc/ha.d/authkeys bravo:/etc/ha.d/
```

Edit the ***ha.cf*** file:

```
vi /etc/ha.d/ha.cf
```

```
debugfile /var/log/ha-debug
 logfile /var/log/ha-log
 logfacility     local0
 keepalive 1
 deadtime 10
 warntime 5
 initdead 60
 udpport 694
 ucast eth0 192.168.1.10
 ucast eth0 192.168.1.20
```

```
auto_failback off
node alpha
node bravo
```

Copy the ***ha.cf*** file to *bravo*:

```
scp /etc/ha.d/ha.cf bravo:/etc/ha.d/
```

Edit the ***haresources*** file, the IP created here will be the IP that our third node refers to.

```
vi /etc/ha.d/haresources
```

```
alpha IPaddr::192.168.5.2/24/eth0
```

Copy the ***haresources*** file to *bravo*:

```
scp /etc/ha.d/haresources bravo:/etc/ha.d/
```

Start the heartbeat service on both servers to bring up the virtual IP:

```
alpha:/# /etc/init.d/heartbeat start
```

```
bravo:/# /etc/init.d/heartbeat start
```

Heartbeat will bring up the new interface (*eth0:0*).

Note: It may take heartbeat up to one minute to bring the interface up.

```
alpha:/# ifconfig eth0:0
```

```
  eth0:0    Link encap:Ethernet  HWaddr 00:08:C7:DB:01:CC
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
```

Section 3: DRBD Configuration

Configuration for DRBD is done via the ***drbd.conf*** file. This needs to be the same on all nodes (*alpha*, *bravo*, *foxtrot*). Please note that the *usage-count* is set to *yes*, which means it will notify Linbit that you have installed DRBD. No personal information is collected. Please see ***this page*** for more information :

```
global { usage-count yes; }

resource data-lower {
  protocol C;
  net {
    shared-secret "LINBIT";
  }
syncer {
 rate 12M;
}

  on alpha {
    device     /dev/drbd1;
    disk       /dev/hdb1;
    address    172.16.6.10:7788;
    meta-disk  internal;
```



```
  }

  on bravo {
    device    /dev/drbd1;
    disk      /dev/hdd1;
    address   172.16.6.20:7788;
    meta-disk internal;
  }
}

  resource data-upper {
  protocol A;
  syncer {
    after data-lower;
    rate 12M;
    al-extents 513;
  }
  net {
       shared-secret "LINBIT";
  }
  stacked-on-top-of data-lower {
       device /dev/drbd3;
       address 192.168.5.2:7788; # IP provided by Heartbeat
       }

  on foxtrot {
    device     /dev/drbd3;
    disk       /dev/sdb1;
    address    192.168.5.3:7788; # Public IP of the backup node
    meta-disk  internal;
  }
}
```

Section 4: Preparing The DRBD Devices

Now that the configuration is in place, create the metadata on *alpha* and *bravo*.

```
alpha:/usr/src/drbd-8.3.0# drbdadm create-md data-lower
```

```
  Writing meta data...
  initializing activity log
  NOT initialized bitmap
New drbd meta data block successfully created.
```

```
bravo:/usr/src/drbd-8.3.0# drbdadm create-md data-lower
```

```
  Writing meta data...
  initialising activity log
  NOT initialized bitmap
New drbd meta data block successfully created.
```

Now start DRBD on *alpha* and *bravo*:

```
alpha:/usr/src/drbd-8.3.0# /etc/init.d/drbd start
```

```
bravo:/usr/src/drbd-8.3.0# /etc/init.d/drbd start
```

Verify that the lower level DRBD devices are connected:

```
```

```
cat /proc/drbd
```

```
  version: 8.3.0 (api:88/proto:86-89)
  GIT-hash: 9ba8b93e24d842f0dd3fb1f9b90e8348ddb95829 build by root@alpha, 2009-02-05 10:36:11
  0: cs:Connected ro:Secondary/Secondary ds:Inconsistent/Inconsistent C r---
ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:19530844
```

Tell *alpha* to become the primary node:

NOTE: As the command states, this is going to overwrite any data on bravo: Now is a good time to go and grab your favorite drink.

```
alpha:/# drbdadm -- --overwrite-data-of-peer primary data-lower

alpha:/# cat /proc/drbd
```

```
  version: 8.3.0 (api:88/proto:86-89)
  GIT-hash: 9ba8b93e24d842f0dd3fb1f9b90e8348ddb95829 build by root@alpha, 2009-02-05 10:36:11
  0: cs:SyncSource ro:Primary/Secondary ds:UpToDate/Inconsistent C r---
  ns:3088464 nr:0 dw:0 dr:3089408 al:0 bm:188 lo:23 pe:6 ua:53 ap:0 ep:1 wo:b oos:16442556
  [==>.................] sync'ed: 15.9% (16057/19073)M
finish: 0:16:30 speed: 16,512 (8,276) K/sec
```

After the data sync has finished, create the meta-data on *data-upper* on *alpha*, followed by *foxtrot*.

Note the resource is data-upper and the *--stacked* option is on alpha only.

```

alpha:~# drbdadm --stacked create-md data-upper
```

```
  Writing meta data...
```

```
  initialising activity log
  NOT initialized bitmap
  New drbd meta data block successfully created.
success
```

```
foxtrot:/usr/src/drbd-8.3.0# drbdadm create-md data-upper
```

```
  Writing meta data...
  initialising activity log
  NOT initialized bitmap
New drbd meta data block sucessfully created.
```

Bring up the stacked resource, then make *alpha* the primary of *data-upper*:

```
alpha:/# drbdadm --stacked adjust data-upper
```

```
foxtrot:~# drbdadm adjust data-upper

foxtrot:~# cat /proc/drbd
```

```
  version: 8.3.0 (api:88/proto:86-89)
  GIT-hash: 9ba8b93e24d842f0dd3fb1f9b90e8348ddb95829 build by root@foxtrot, 2009-02-02 10:28:37
  1: cs:Connected ro:Secondary/Secondary ds:Inconsistent/Inconsistent A r---
ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:19530208
```

```
alpha:~# drbdadm --stacked -- --overwrite-data-of-peer primary data-upper

alpha:~#  cat /proc/drbd
```

```
  version: 8.3.0 (api:88/proto:86-89)
  GIT-hash: 9ba8b93e24d842f0dd3fb1f9b90e8348ddb95829 build by root@alpha, 2009-02-05 10:36:11
  0: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r---
  ns:19532532 nr:0 dw:1688 dr:34046020 al:1 bm:1196 lo:156 pe:0 ua:0 ap:156 ep:1 wo:b oos:0
  1: cs:SyncSource ro:Primary/Secondary ds:UpToDate/Inconsistent A r---
  ns:14512132 nr:0 dw:0 dr:14512676 al:0 bm:885 lo:156 pe:32 ua:292 ap:0 ep:1 wo:b oos:5018200
  [==============>......] sync'ed: 74.4% (4900/19072)M
finish: 0:07:06 speed: 11,776 (10,992) K/sec
```

Drink time again!

After the sync is complete, access your DRBD block device via `/dev/drbd3`. This will write to both local nodes and the remote third node. In your Heartbeat configuration you will use the "drbdupper" script to bring up your `/dev/drbd3` device. Have fun!

DRBD® and LINBIT® are registered trademarks of LINBIT, Austria.