

## **HA CLUSTER**

To look into:

### **HA**

1. Add crm and gui
2. More testing of current nodes
3. Add more nodes

### **Application Server**

1. Open Source: Jboss (Java), glassfish (Java), Zope (Python)
2. Commercial: Oracle, Sun, BEA, IBM <-- all java too.

### **Applications to try**

1. Zimbra (javascript, AJAX)

### **Storage Server**

1. Openfiler

Other things to look into:

<http://mln.sourceforge.net/>

Venti for guest image archival purposes

-----  
Install Xen3.

Using Edgy's version of Xen caused problems with hyperthreading.

There is a known bug on this <https://launchpad.net/distros/ubuntu/+source/xen-source-2.6.17/+bug/71348>

We could only get one guest running at a time. A second guest would cause a kernel panic. Disabling hyperthreading allowed two simultaneous guests (the third causing a crash). We decided to download the latest stable Xen kernel and use that (the Xen 3.0.3 kernel allows us to use more guest with no crashes) we downloaded from: [http://bits.xensource.com/oss-xen/release/3.0.3-0/bin.tgz/xen-3.0.3\\_0-install-x86\\_32.tgz](http://bits.xensource.com/oss-xen/release/3.0.3-0/bin.tgz/xen-3.0.3_0-install-x86_32.tgz)

Create a guest and install heartbeat-2 (version 2.0.7).

### **Fixes for heartbeat:**

needed to make `/var/run/heartbeat` before starting heartbeat in the init script. (?)

Copy the guest for each node of the cluster and change config files:

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

`/etc/ha.d/nodeinfo` <-- this file contains the nodename

Customize hostname files: (`/etc/hosts`, `/etc/hostname` etc.)

### **Initial setup**

2 nodes using bcast (to broadcast their status)

adding more nodes using bcast didn't work.

### **Setup as of 12/13/06**

Next, made two sets of 2 nodes using two different multicast addresses and two different udp ports

the ha.cf config file for each of the nodes:

#### **Node1:**

```
node node1
node node2
debugfile /var/log/ha-debug
logfile /var/log/ha-log
logfacility local0
deadtime 5
udpport 794
mcast eth0 225.1.1.1 794 1 0
auto_failback off
```

#### **Node2:**

```
node node1
node node2
debugfile /var/log/ha-debug
logfile /var/log/ha-log
logfacility local0
deadtime 5
udpport 794
mcast eth0 225.1.1.1 794 1 0
auto_failback off
```

#### **Node3:**

```
node node3
node node4
debugfile /var/log/ha-debug
logfile /var/log/ha-log
logfacility local0
deadtime 5
udpport 694
mcast eth0 225.220.220.220 694 1 0
auto_failback off
```

#### **Node 4:**

```
node node4
node node3
debugfile /var/log/ha-debug
logfile /var/log/ha-log
logfacility local0
deadtime 5
udpport 694
mcast eth0 225.220.220.220 694 1 0
auto_failback off
```

#### **openfiler xen config**

```
kernel = "/boot/vmlinuz-2.6.16.33-xen"
ramdisk = "/boot/initrd.img-2.6.16.33-xen"
disk = ['tap:aio:/xen/images/openfiler-2.2.img,xvda,w','tap:aio:/xen/images/openfiler-
volumes/volume1.img,xvdb,w']
root = "/dev/xvda1 ro"
vif = ['mac=00:16:3e:6d:6e:7a']
```

name = "openfiler"

## Problems and solutions

When copying and renaming nodes in the cluster, heartbeat gave us a lot of warning about node names uid's changing.

To fix this we cleared out the two files in /var/lib/heartbeat **hb\_generation** and **hb\_uuid**

For this we had to look through the code to see how uuid generation was done. The uuid are stored in a hash table, so

we found the files about that it stored the hash tables in binary files.

Downloading the heartbeat source:

install **mercurial**

**hg clone <http://hg.linux-ha.org/dev>**

Useful also to look on [google.com/codesearch](http://google.com/codesearch)

Running more than two nodes on the same broadcast caused errors and weren't detecting the status of the other nodes.

This is the reason behind our switch to mcast (**multicast**), which allowed us to have multiple pairs of nodes running at once.

Problems with the Xen version of **Openfiler** working correctly. It couldn't find any volumes to use as storage space, or

including the volume that openfiler was stored on. This is fixed by creating them as **xvda** volumes instead of **hda** or **sda**.

[https://issues.rpath.com/browse/RPL-814#action\\_22344](https://issues.rpath.com/browse/RPL-814#action_22344)

[http://wiki.rpath.com/wiki/Xen\\_DomU\\_Guide](http://wiki.rpath.com/wiki/Xen_DomU_Guide)

<http://www.openfiler.com/community/forums/viewtopic.php?p=1849#1849>

Not a lot of xen documentation on xvd, this should be fixed.

This is what Sean said about xvd:

"xvd is the preferred way to specify block devices, as the xen vbd interface doesn't truly emulate either IDE or SCSI, and allows the Xen guest to understand that it has virtual disks instead of native hardware"

hack needed in our particular case, since the network card wasn't showing up and or working soon enough for heartbeat to load. The fix was to add a sleep 5 in the start\_heartbeat function init.d script (start case)

openfiler hacks needed:

<http://i-hug.sarovar.org/wiki/pmwiki.php/InstallGuide/QemuCrossHurd>

## Clarkson specific

Netreg the mac addresses for the xen guests