

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

**LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.**

www.theredhatsummit.com

FOLLOW US ON TWITTER

www.twitter.com/redhatsummit

TWEET ABOUT IT

[#summitjbw](https://twitter.com/summitjbw)

READ THE BLOG

<http://summitblog.redhat.com/>

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Clustered Applications and Filesystems with Red Hat Enterprise Linux 6

Thomas Cameron, RHCA
Managing Solutions Architect, Red Hat
Lon Hohberger
Principal Software Engineer, Red Hat
June 23rd, 2010

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Agenda

Red Hat and Virtualization

Architecture

Configure Raw Storage and iSCSI

- On the Target
- On the Client

Define the Cluster

Define Fencing

Define Clustered Resources

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Agenda

Define a Clustered Service

Test Failover

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Red Hat and Clustering

Red Hat leads the way in Open Source clustering

Acquired Sistina for \$31 million in early 2004, including Global Filesystem and ClusterSuite.

Made the code Open Source in mid-2004.

GFS and ClusterSuite formerly separate, layered products from the base distribution.

With RHEL 5, included in Advanced Platform.

RHEL now includes clustered filesystem (GFS2), clustered LVM, multipath IO and high availability application services with the distribution.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



DISCLAIMER

Today's presentation was developed on a pre-release, beta version of Red Hat Enterprise Linux version 6. This presentation is not a commitment by Red Hat to include any of the technologies presented today in any future release of Red Hat Enterprise Linux.

If you try any of this at home and it breaks, we only guarantee that you get to keep all the pieces. Assuming they don't burn up.

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Architecture

Two node cluster - case.tc.redhat.com and molly.tc.redhat.com - each installed with @base and @web-server

case and molly are managed by neuromancer.tc.redhat.com - installed with @base

neuromancer is also sharing storage via iSCSI

Fence device is a WTI managed power distribution unit

case and molly will be running Apache as a clustered application, using GFS2 as back-end storage.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Architecture

As this is beta software, and for the sake of brevity and clarity, SELinux and iptables are disabled.


At GA you should absolutely use iptables and SELinux!

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT





SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Configure the iSCSI Target on neuromancer

Install scsi-target-utils, part of the Network Storage Server group, on neuromancer

“yum groupinstall storage-server” will install:

- scsi-target-utils
- perl-Config-General
- librdmacm
- libibverbs
- libmlx4



```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# yum -y groupinstall storage-server  
Loaded plugins: rhnplugin  
This system is not registered with RHN.  
RHN support will be disabled.  
Setting up Group Process  
Resolving Dependencies  
--> Running transaction check  
---> Package scsi-target-utils.x86_64 0:1.0.4-0.el6 set to be updated  
--> Processing Dependency: perl(Config::General) for package: scsi-target-utils-1.0.4-0.el6.x86_64  
--> Processing Dependency: libibverbs.so.1()(64bit) for package: scsi-target-utils-1.0.4-0.el6.x86_64  
--> Processing Dependency: librdmacm.so.1()(64bit) for package: scsi-target-utils-1.0.4-0.el6.x86_64  
--> Running transaction check  
---> Package libibverbs.x86_64 0:1.1.3-4.el6 set to be updated  
--> Processing Dependency: libibverbs-driver for package: libibverbs-1.1.3-4.el6.x86_64  
---> Package librdmacm.x86_64 0:1.0.10-2.el6 set to be updated  
---> Package perl-Config-General.noarch 0:2.44-1.el6 set to be updated  
--> Running transaction check  
---> Package libmlx4.x86_64 0:1.0.1-5.el6 set to be updated  
--> Finished Dependency Resolution
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Prepare shared storage on neuromancer

For demonstration purposes, /dev/sda4 is created and will be shared via iSCSI.

Use fdisk to prepare the partition.

- fdisk /dev/sda and create the partition
- Create a new primary partition and define the start and save your changes
- Reboot or run partprobe
- Don't make any filesystem on this partition



```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# fdisk /dev/sda  
  
WARNING: DOS-compatible mode is deprecated. It's strongly recommended to  
switch off the mode (command 'c') and change display units to  
sectors (command 'u').  
  
Command (m for help): n  
Command action  
  e   extended  
  p   primary partition (1-4)  
p  
Selected partition 4  
First cylinder (3225-12188, default 3225):  
Using default value 3225  
Last cylinder, +cylinders or +size{K,M,G} (3225-12188, default 12188):  
Using default value 12188  
  
Command (m for help): w  
The partition table has been altered!  
  
Calling ioctl() to re-read partition table.  
  
WARNING: Re-reading the partition table failed with error 16: Device or resource  
busy.  
The kernel still uses the old table. The new table will be used at  
the next reboot or after you run partprobe(8) or kpartx(8)  
Syncing disks.  
[root@neuromancer ~]# █
```



Prepare shared storage on neuromancer

Edit the iSCSI target configuration file,
`/etc/tgt/targets.conf`

Add the target and define the initiators for security purposes.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@neuromancer:~  
File Edit View Search Terminal Help  
  
# Continue if tgtadm exits with non-zero code (equivalent of  
# --ignore-errors command line option)  
#ignore-errors yes  
  
# Sample target with one LUN only. Defaults to allow access for all initiators:  
  
#<target iqn.2008-09.com.example:server.target1>  
#   backing-store /dev/LVM/somedevice  
#</target>  
  
<target iqn.2010-06-com.redhat.tc:neuromancer.target1>  
    backing-store /dev/sda4  
    initiator-address 172.31.100.3  
    initiator-address 172.31.100.4  
</target>  
  
# Similar, but we use "direct-store" instead of "backing-store".  
# "direct-store" reads drive parameters with sg_inq command and sets them to  
# the target.  
# Parameters fatched with sg_inq are:  
-- INSERT --
```



Prepare shared storage on neuromancer

chkconfig tgtd on and service tgtd restart

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# vi /etc/tgt/targets.conf  
[root@neuromancer ~]# vi /etc/tgt/targets.conf  
[root@neuromancer ~]# chkconfig tgtd on  
[root@neuromancer ~]# service tgtd restart  
Stopping SCSI target daemon: not running [ FAILED ]  
Starting SCSI target daemon: [ OK ]  
[root@neuromancer ~]#
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Prepare shared storage on neuromancer

Verify by running “tgtadm --lld iscsi --mode target --op show”

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# tgtadm --lld iscsi --mode target --op show  
Target 1: iqn.2010-06-com.redhat.tc:neuromancer.target1  
  System information:  
    Driver: iscsi  
    State: ready  
  I_T nexus information:  
  LUN information:  
    LUN: 0  
      Type: controller  
      SCSI ID: IET      00010000  
      SCSI SN: beaf10  
      Size: 0 MB  
      Online: Yes  
      Removable media: No  
      Backing store type: rdwr  
      Backing store path: None  
    LUN: 1  
      Type: disk  
      SCSI ID: IET      00010001  
      SCSI SN: beaf11  
      Size: 73724 MB  
      Online: Yes  
      Removable media: No  
      Backing store type: rdwr  
      Backing store path: /dev/sda4  
  Account information:  
  ACL information:  
    172.31.100.3  
    172.31.100.4  
[root@neuromancer ~]#
```



Connect to shared storage on both clients

Install iscsi-initiator-utils, which is in the iSCSI Storage Client group, on case and molly.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@case:~  
File Edit View Search Terminal Help  
[root@case ~]# yum -y groupinstall storage-client-iscsi  
Loaded plugins: rhnplugin  
This system is not registered with RHN.  
RHN support will be disabled.  
Setting up Group Process  
Resolving Dependencies  
--> Running transaction check  
---> Package iscsi-initiator-utils.x86_64 0:6.2.0.872-4.el6 set to be updated  
--> Finished Dependency Resolution  
  
Dependencies Resolved  
  
=====
```

Package	Arch	Version	Repository	Size
---------	------	---------	------------	------

```
=====
```

Installing:

iscsi-initiator-utils	x86_64	6.2.0.872-4.el6	Server	872 k
-----------------------	--------	-----------------	--------	-------

```
=====
```

Transaction Summary

```
=====
```

Install	1 Package(s)
Upgrade	0 Package(s)

```
=====
```

Total download size: 872 k
Installed size: 2.6 M
Downloading Packages:
iscsi-initiator-utils-6.2.0.872-4.el6.x86_64.rpm | 872 kB 00:00
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
Installing : iscsi-initiator-utils-6.2.0.872-4.el6.x86_64 1/1

Installed:
iscsi-initiator-utils.x86_64 0:6.2.0.872-4.el6

Complete!
[root@case ~]#

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Connect to shared storage on both clients

chkconfig on and start the iscsid and iscsi services

Use iscsiadm to query the target:

- iscsiadm --mode discovery --type sendtargets --portal 172.31.100.14
- Use the ip address, not the domain name



```
root@case:~  
File Edit View Search Terminal Help  
[root@case ~]# service iscsid start  
Starting iscsid: [ OK ]  
[root@case ~]# service iscsi start  
Starting iscsi: [ OK ]  
[root@case ~]# iscsiadm --mode discovery --type sendtargets --portal 172.31.100.14  
172.31.100.14:3260,1 iqn.2010-06-com.redhat.tc:neuromancer.target1  
[root@case ~]#
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Connect to shared storage on both clients

On both servers, log into the target:

- `iscsiadm --mode node --targetname iqn.2010-06-com.redhat.tc:neuromancer.target1 --portal 172.31.100.14 --login`

You should see new storage available to the server. Test by looking at `/proc/partitions`



```
root@case:~
File Edit View Search Terminal Help
[root@case ~]# cat /proc/partitions
major minor #blocks name

 8          0 488386584 sda
 8          1   204800 sda1
 8          2 215040000 sda2
 8          3  4194304 sda3
 8         16 488386584 sdb
[root@case ~]# iscsiadm --mode node --targetname iqn.2010-06-com.redhat.tc:neuromancer.target1 --portal 172.31.100.14 --login
Logging in to [iface: default, target: iqn.2010-06-com.redhat.tc:neuromancer.target1, portal: 172.31.100.14,3260]
Login to [iface: default, target: iqn.2010-06-com.redhat.tc:neuromancer.target1, portal: 172.31.100.14,3260] successful.
[root@case ~]# cat /proc/partitions
major minor #blocks name

 8          0 488386584 sda
 8          1   204800 sda1
 8          2 215040000 sda2
 8          3  4194304 sda3
 8         16 488386584 sdb
 8         32 71995982 sdc
[root@case ~]#
```



Connect to shared storage on both clients

Don't worry about making a filesystem yet, we'll do that once the cluster software is installed and the cluster is running

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Install the clustering software on the nodes

There are two components of the cluster

- High availability application service
- Clustered filesystem (GFS2)



Install the clustering software on the nodes

Install the high availability software first

- The yum group is High Availability

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT





Computer



root's Home



Trash



500GB

```
root@case:~  
File Edit View Search Terminal Help  
Installing : openaislib-1.1.1-5.el6.x86_64 15/45  
Installing : modcluster-0.16.2-7.el6.x86_64 16/45  
Installing : 1:telnet-0.17.46-1.el6.x86_64 17/45  
Installing : yajl-1.3.2-1.el6.x86_64  
Installing : fence-1:1.10.1-1.el6.x86_64  
Installing : nss-tls-1:3.44-1.el6.x86_64  
Installing : ricci-0.15.0-4.el6.x86_64  
Installing : ipmitool-1.8.16-1.el6.x86_64  
Installing : lm_sensors-3.4.2-1.el6.x86_64  
Installing : 1:net-tools-2.0-22.el6.x86_64  
Installing : 1:net-tools-2.0-22.el6.x86_64  
Installing : libevent-2.0.5-2.el6.x86_64  
Installing : nfs-utils-1:1.24.0-1.el6.x86_64  
Installing : 1:nfs-utils-1:1.24.0-1.el6.x86_64  
Installing : libtasn1-0.4.8-3.el6.x86_64  
Installing : cyrus-sasl-2.1.23-15.el6.x86_64  
Installing : keyutils-1:1.6-13.el6.x86_64  
Installing : cifs-utils-2.3.1-1.el6.x86_64  
Installing : nc-1:1.89-1.el6.x86_64  
Installing : sg3_utils-1:1.36-1.el6.x86_64  
Installing : sg3_utils-1:1.36-1.el6.x86_64  
Installing : numactl-1:2.0.9-4.el6.x86_64  
Installing : libvirt-1:1.11.0-4.el6.x86_64
```

```
root@molly:~  
File Edit View Search Terminal Help  
[root@molly ~]# yum -y groupinstall ha  
Loaded plugins: rhnplugin  
This system is not registered with RHN.  
RHN support will be disabled.  
Setting up Group Process  
Resolving Dependencies  
--> Running transaction check  
---> Package cman.x86_64 0:3.0.12-6.el6 set to be updated  
--> Processing Dependency: clusterlib = 3.0.12-6.el6 for package: cman-3.0.12-6.el6.x86_64  
--> Processing Dependency: fence-virt >= 0.2.1-1 for package: cman-3.0.12-6.el6.x86_64  
--> Processing Dependency: modcluster >= 0.15.0-3 for package: cman-3.0.12-6.el6.x86_64  
--> Processing Dependency: openais >= 1.1.1-1 for package: cman-3.0.12-6.el6.x86_64  
--> Processing Dependency: ricci >= 0.15.0-4 for package: cman-3.0.12-6.el6.x86_64  
--> Processing Dependency: corosync >= 1.2.2-1 for package: cman-3.0.12-6.el6.x86_64  
--> Processing Dependency: fence-agents >= 3.0.6-1 for package: cman-3.0.12-6.el6.x86_64  
--> Processing Dependency: libcpng.so.4(COROSYNC_CPG_1.0)(64bit) for package: cman-3.0.12-6.el6.x86_64
```

Install the clustering software on the nodes

Install the filesystem clustering software next

- The yum group is Clustered Filesystem

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





root's Home



Trash



500GB



Screen

```
root@case:~  
File Edit View Search Terminal Help  
Transaction Summary  
-----  
Install      2 Package(s)  
Upgrade      0 Package(s)  
  
Total download size: 521  
Installed size: 1.2 M  
Downloading Packages:  
(1/2): gfs2-utils-3.0.12  
(2/2): lvm2-cluster-2.02  
-----  
Total  
Running rpm_check_debug  
Running Transaction Test  
Transaction Test Succeeded  
Running Transaction  
  Installing      : gfs2-  
  Installing      : lvm2-  
  
Installed:  
  gfs2-utils.x86_64 0:3.0.12-6.el6  
  
Complete!  
[root@case ~]#
```

```
root@molly:~  
File Edit View Search Terminal Help  
[root@molly ~]# yum -y groupinstall clustered-storage  
Loaded plugins: rhnplugin  
This system is not registered with RHN.  
RHN support will be disabled.  
Setting up Group Process  
Resolving Dependencies  
--> Running transaction check  
---> Package gfs2-utils.x86_64 0:3.0.12-6.el6 set to be updated  
---> Package lvm2-cluster.x86_64 0:2.02.66-3.el6 set to be updated
```

Install the clustering software on the nodes

`chkconfig ricci on`

`service ricci start (or reboot)`

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



```
root@tcx60:~/Desktop
File Edit View Search Terminal Help
Installed size: 1.2 M
Downloading Packages:
(1/2): gfs2-utils-3.0.12-6.el6.x86_64.rpm | 235 kB 00:00
(2/2): lvm2-cluster-2.02.66-3.el6.x86_64.rpm | 286 kB 00:00
-----
Total 30 MB/s | 521 kB 00:00
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing      : gfs2-utils-3.0.12-6.el6.x86_64 1/2
  Installing      : lvm2-cluster-2.02.66-3.el6.x86_64 2/2

Installed:
  gfs2-utils.x86_64 0:3.0.12-6.el6      lvm2-cluster.x86_64 0:2.02.66-3.el6

Complete!
[root@molly ~]# chkconfig ricci --list
ricci      0:off  1:off  2:off  3:off  4:off  5:off  6:off
[root@molly ~]# chkconfig ricci on
[root@molly ~]# init 6; exit
logout
Connection to molly closed.
[root@tcx60 Desktop]#
```



Install the cluster management software on the management server

The yum group is “High Availability Management”

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# yum -y groupinstall ha-management  
Loaded plugins: rhnplugin  
This system is not registered with RHN.  
RHN support will be disabled.  
Setting up Group Process  
Resolving Dependencies  
--> Running transaction check  
---> Package luci.x86_64 0:0.22.2-2.el6 set to be updated  
--> Processing Dependency: python-paste >= 1.7.2-5.el6 for package: luci-0.22.2-2.el6.x86_64  
--> Processing Dependency: TurboGears2 for package: luci-0.22.2-2.el6.x86_64  
--> Processing Dependency: python-tw-forms for package: luci-0.22.2-2.el6.x86_64  
--> Processing Dependency: python-repoze-who-friendlyform for package: luci-0.22.2-2.el6.x86_64  
--> Running transaction check  
---> Package TurboGears2.noarch 0:2.0.3-4.el6 set to be updated  
--> Processing Dependency: python-turbojson >= 1.2.1 for package: TurboGears2-2.0.3-4.el6.noarch  
--> Processing Dependency: python-toscawidgets >= 0.9.4 for package: TurboGears2-2.0.3-4.el6.noarch  
--> Processing Dependency: python-weberror >= 0.10.1 for package: TurboGears2-2.0.3-4.el6.noarch  
--> Processing Dependency: python-genshi >= 0.5.1 for package: TurboGears2-2.0.3-4.el6.noarch
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Install the cluster management software on the management server

`chkconfig luci on`

`service luci start`

Open the URL listed when luci starts

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



```
root@neuromancer:~  
File Edit View Search Terminal Help  
python-repoze-who-testutil.noarch 0:1.0-0.4.rc1.el6  
python-routes.noarch 0:1.10.3-2.el6  
python-setuptools.noarch 0:0.6.10-1.el6  
python-simplejson.x86_64 0:2.0.9-3.1.el6  
python-sqlalchemy.noarch 0:0.5.5-2.1.el6  
python-tempita.noarch 0:0.4-2.el6  
python-toscawidgets.noarch 0:0.9.8-1.el6  
python-transaction.noarch 0:1.0-0.5.a1.el6  
python-turbojson.noarch 0:1.2.1-8.1.el6  
python-tw-forms.noarch 0:0.9.8-1.el6  
python-weberror.noarch 0:0.10.1-4.el6  
python-webflash.noarch 0:0.1-0.2.a9.el6  
python-webhelpers.noarch 0:0.6.4-4.el6  
python-webob.noarch 0:0.9.6.1-3.el6  
python-webtest.noarch 0:1.2-2.el6  
python-zope-filesystem.x86_64 0:1-4.1.el6  
python-zope-interface.x86_64 0:3.5.2-2.1.el6  
python-zope-sqlalchemy.noarch 0:0.4-2.el6  
  
Complete!  
[root@neuromancer ~]# chkconfig luci on  
[root@neuromancer ~]# service luci restart  
Point your web browser to https://neuromancer.tc.redhat.com:8084 to access luci  
[root@neuromancer ~]#
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Install the cluster management software on the management server

As the Conga web application uses a self-signed SSL certificate, you'll get a warning. This is normal and expected. Accept the certificate and continue.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





This Connection is Untrusted

You have asked Firefox to connect securely to **neuromancer.tc.redhat.com:8084**, but we can't confirm that your connection is secure.

Normally, when you try to connect securely, sites will present trusted identification to prove that you are going to the right place. However, this site's identity can't be verified.

What Should I Do?

If you usually connect to this site without problems, this error could mean that someone is trying to impersonate the site, and you shouldn't continue.

Get me out of here!

- ▶ **Technical Details**
- ▶ **I Understand the Risks**

Install the cluster management software on the management server

Log into the web UI as root

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT





File Edit View History Bookmarks Tools Help

Navigation icons: back, forward, stop, reload, home, search, redhat.com icon

Address bar: [https://neuromancer.tc.redhat.com:8084/login?came_from=https'](https://neuromancer.tc.redhat.com:8084/login?came_from=https)

Search bar: Google

Most Visited Red Hat Red Hat Magazine Red Hat Network Red Hat Support

Login Form

[About](#) [Login](#)

Homebase

Login

Username

Password



Create the cluster

From the main Conga page, choose “Manage Clusters”

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





Homebase

Homebase

Manage Clusters

NOTIFICATIONS

CLUSTER SUMMARY

Name	Status	Nodes Joined
------	--------	--------------



Create the cluster

Choose “Create”

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Create the cluster

Fill out the cluster name

“Use the same passwords for all nodes” is not working (BZ599074).

Fill out node by FQDN or ip and enter root's password

Choose download packages to ensure your servers grab the latest versions from RHN

Choose “Reboot nodes...” (not working, BZ599080)

Choose “Enable shared storage...”


Commit your changes and wait up to 30 seconds

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





Create a Cluster

Cluster Name summit-cluster

☐ Use same password for all nodes

Node Hostname

Password

Node Port

molly.tc.redhat.com

●●●●●●●●●●

11111



case.tc.redhat.com

●●●●●●●●●●

11111




Add Another Node

☒ Download Packages☐ Use locally installed packages☒ Reboot nodes before joining cluster☒ Enable shared storage support

Create Cluster

Cancel



i Creating the cluster "summit-cluster", this may take up to 30 seconds. Please refresh screen to view created cluster.

Homebase

Manage Clusters

+ Add ★ Create ✕ Delete					
Name	Status	Possible Votes	Current Votes	Quorum	Nodes Joined

Create the cluster

Click on “Manage Clusters”

If your cluster shows up with a green light next to it, you're in business!

If it doesn't, typically one of a few things is wrong:

- name resolution
- firewall rules
- selinux (until GA)


Reboot the cluster manually to make sure it comes up correctly.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





Create the cluster

I also like to see check that the cluster processes started up.

You can use clustat after the cluster comes up to check cluster status.



```
root@case:~  
File Edit View Search Terminal Help  
2096 ?      SLsl  0:00 clvmd -T30  
2100 ?      S      0:00 [dlm_astd]  
2101 ?      S      0:00 [dlm_scand]  
2102 ?      S      0:00 [dlm_recv/0]  
2103 ?      S      0:00 [dlm_recv/1]  
2104 ?      S      0:00 [dlm_recv/2]  
2105 ?      S      0:00 [dlm_recv/3]  
2106 ?      S      0:00 [dlm_send]  
2107 ?      S      0:00 [dlm_recoverd]  
2132 ?      S<Ls  0:00 rgmanager  
2134 ?      S<l   0:00 rgmanager  
2136 ?      S      0:00 [dlm_recoverd]  
5026 pts/0  R+    0:00 ps ax  
^C  
[root@case ~]# clustat  
Cluster Status for summit-cluster @ Sat Jun 19 21:00:37 2010  
Member Status: Quorate  
  
Member Name              ID   Status  
-----  
molly.tc.redhat.com      1   Online  
case.tc.redhat.com       2   Online, Local  
[root@case ~]#
```

```
2024 ?      S      0:00 [dlm_astd]  
2025 ?      S      0:00 [dlm_scand]  
2026 ?      S      0:00 [dlm_recv/0]  
2027 ?      S      0:00 [dlm_recv/1]  
2028 ?      S      0:00 [dlm_send]  
2029 ?      S      0:00 [dlm_recoverd]  
2056 ?      S<Ls  0:00 rgmanager  
2058 ?      S<l   0:00 rgmanager  
2061 ?      S      0:00 [dlm_recoverd]  
5212 ?      R<    0:00 /usr/sbin/clustat -f -x  
5217 pts/0  R+    0:00 ps ax
```

tty5

tty6

Create Shared Storage

Now that the cluster is up, we can set up the shared storage from the hosts.

Enable each node to be able to use clustered logical volumes:

- `lvmconf --enable-cluster`



```
root@molly:~  
File Edit View Search Terminal Help  
[root@molly ~]# grep locking_type /etc/lvm/lvm.conf  
    locking_type = 1  
    # NB. This option only affects locking_type = 1 viz. local file-based  
    # The external locking library to load if locking_type is set to 2.  
[root@molly ~]# lvmconf --enable-cluster  
[root@molly ~]# grep locking_type /etc/lvm/lvm.conf  
    locking_type = 3  
    # NB. This option only affects locking_type = 1 viz. local file-based  
    # The external locking library to load if locking_type is set to 2.  
[root@molly ~]#
```



Create Shared Storage

Create a partition on the shared storage we connected to via iSCSI

- fdisk the disk
 - new partition
 - primary
 - start at the first cylinder
 - end at the last
 - write changes



```
root@molly:~  
File Edit View Search Terminal Help  
[root@molly ~]# fdisk /dev/sdb  
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel  
Building a new DOS disklabel with disk identifier 0xdaf27d38.  
Changes will remain in memory only, until you decide to write them.  
After that, of course, the previous content won't be recoverable.  
  
Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)  
  
WARNING: DOS-compatible mode is deprecated. It's strongly recommended to  
switch off the mode (command 'c') and change display units to  
sectors (command 'u').  
  
Command (m for help): n  
Command action  
   e   extended  
   p   primary partition (1-4)  
p  
Partition number (1-4): 1  
First cylinder (1-8963, default 1):  
Using default value 1  
Last cylinder, +cylinders or +size{K,M,G} (1-8963, default 8963):  
Using default value 8963  
  
Command (m for help): w  
The partition table has been altered!  
  
Calling ioctl() to re-read partition table.  
Syncing disks.  
[root@molly ~]#
```



Create Shared Storage

Verify the other node sees the new partition

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@case:~  
File Edit View Search Terminal Help  
[root@case ~]# fdisk -l /dev/sdc  
  
Disk /dev/sdc: 73.7 GB, 73723885568 bytes  
255 heads, 63 sectors/track, 8963 cylinders  
Units = cylinders of 16065 * 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0xdaf27d38  
  
   Device Boot      Start         End      Blocks   Id  System  
/dev/sdc1             1         8963     71995266   83   Linux  
[root@case ~]#
```



Create Shared Storage

Create a logical volume on the shared disk

- `pvcreate /dev/sdb1`
- `vgcreate -c y ApacheVol /dev/sdb1`
- `lvcreate -L 10G -n ApacheLV ApacheVol`



```
root@molly:~  
File Edit View Search Terminal Help  
[root@molly ~]# pvcreate /dev/sdb1  
Physical volume "/dev/sdb1" successfully created  
[root@molly ~]# vgcreate -c y ApacheVol /dev/sdb1  
Clustered volume group "ApacheVol" successfully created  
[root@molly ~]# lvcreate -L 10G -n ApacheLV ApacheVol  
Logical volume "ApacheLV" created  
[root@molly ~]#
```



Create Shared Storage

You **might** need to stop and start the clvmd service on the other node.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@case:~  
File Edit View Search Terminal Help  
[root@case ~]# service clvmd stop  
Deactivating clustered VG(s): 0 logical volume(s) in volume group "ApacheVol"  
now active  
[ OK ]  
Signaling clvmd to exit [ OK ]  
clvmd terminated [ OK ]  
[root@case ~]# service clvmd start  
Starting clvmd: [ OK ]  
Activating VG(s): 1 logical volume(s) in volume group "ApacheVol" now active  
[ OK ]  
[root@case ~]#
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Create Shared Storage

Create a GFS2 filesystem on the shared disk

- `mkfs.gfs2 -j 3 -p lock_dlm -t summit-cluster:shared-web /dev/mapper/ApacheVol-ApacheLV`

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@molly:~  
File Edit View Search Terminal Help  
[root@molly ~]# mkfs.gfs2 -j 3 -p lock_dlm -t summit-cluster:shared-web /dev/mapper/ ApacheVol-ApacheLV  
This will destroy any data on /dev/mapper/ ApacheVol-ApacheLV.  
It appears to contain: symbolic link to `../dm-0`  
  
Are you sure you want to proceed? [y/n] y  
  
Device: /dev/mapper/ ApacheVol-ApacheLV  
Blocksize: 4096  
Device Size 10.00 GB (2621440 blocks)  
Filesystem Size: 10.00 GB (2621438 blocks)  
Journals: 3  
Resource Groups: 40  
Locking Protocol: "lock_dlm"  
Lock Table: "summit-cluster:shared-web"  
UUID: E2197468-C7F0-841B-033E-EA5A8C86F8C2  
  
[root@molly ~]#
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Create Shared Storage

Since this is a shared filesystem, we can mount it at the host level via `/etc/fstab`.

- Mount the GFS2 volume from the command line
- `tail -1 /etc/mtab` to see the format for `fstab`




```
root@molly:~  
File Edit View Search Terminal Help  
[root@molly ~]# mount -o noatime /dev/mapper/AdobeVol-ApacheLV /var/www/html/  
[root@molly ~]# df -h  
Filesystem      Size  Used Avail Use% Mounted on  
/dev/sda2       21G  1012M   19G   6% /  
tmpfs           3.9G   28M   3.8G   1% /dev/shm  
/dev/sda1       194M   25M   160M  14% /boot  
/dev/dm-0       10G   388M   9.7G   4% /var/www/html  
[root@molly ~]# tail /var/log/messages  
Jun 19 22:17:07 molly kernel: GFS2: fsid=: Trying to join cluster "lock_dlm", "summit-cluster:shared-web"  
Jun 19 22:17:07 molly kernel: GFS2: fsid=summit-cluster:shared-web.1: Joined cluster. Now mounting FS...  
Jun 19 22:17:07 molly kernel: GFS2: fsid=summit-cluster:shared-web.1: jid=1, already locked for use  
Jun 19 22:17:07 molly kernel: GFS2: fsid=summit-cluster:shared-web.1: jid=1: Looking at journal...  
Jun 19 22:17:07 molly kernel: GFS2: fsid=summit-cluster:shared-web.1: jid=1: Done  
Jun 19 22:19:41 molly kernel: GFS2: fsid=: Trying to join cluster "lock_dlm", "summit-cluster:shared-web"  
Jun 19 22:19:41 molly kernel: GFS2: fsid=summit-cluster:shared-web.1: Joined cluster. Now mounting FS...  
Jun 19 22:19:41 molly kernel: GFS2: fsid=summit-cluster:shared-web.1: jid=1, already locked for use  
Jun 19 22:19:41 molly kernel: GFS2: fsid=summit-cluster:shared-web.1: jid=1: Looking at journal...  
Jun 19 22:19:41 molly kernel: GFS2: fsid=summit-cluster:shared-web.1: jid=1: Done  
[root@molly ~]#
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT





A terminal window titled "root@case:~" with a menu bar containing "File", "Edit", "View", "Search", "Terminal", and "Help". The terminal shows the command `tail -1 /etc/mtab` being executed, resulting in the output `/dev/dm-0 /var/www/html gfs2 rw,noatime,hostdata=jid=0 0 0`. The prompt `[root@case ~]#` is visible at the end of the line.

```
root@case:~  
File Edit View Search Terminal Help  
[root@case ~]# tail -1 /etc/mtab  
/dev/dm-0 /var/www/html gfs2 rw,noatime,hostdata=jid=0 0 0  
[root@case ~]#
```

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@molly:~
File Edit View Search Terminal Help

#
# /etc/fstab
# Created by anaconda on Sat Jun 19 13:49:54 2010
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=4d3090e3-a76c-4eed-a896-0838b4d4c7d7 / ext4 default
ts 1 1
UUID=28bc3a2b-5bda-4f05-ab60-f0e0dbccd2e4 /boot ext4 default
ts 1 2
UUID=72cdd449-8f99-4379-9183-2f6f27bf1d89 swap swap default
ts 0 0
tmpfs /dev/shm tmpfs defaults 0 0
devpts /dev/pts devpts gid=5,mode=620 0 0
sysfs /sys sysfs defaults 0 0
proc /proc proc defaults 0 0
/dev/dm-0 /var/www/html gfs2 rw,noatime 0 0
~
~
~
~
-- INSERT --
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Set up a shared fence device

Fence devices can be any one of a number of mechanisms for blocking a crashed node from data. Some examples include:

- managed power switches from APC and WTI
- HP iLO
- Dell DRAC
- IBM RSA
- Fiber switches
- Virtual machine fencing



Set up a shared fence device


Navigate Manage Clusters, click on the cluster, then click on the Fence Devices tab, then click “Add”

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT






Set up a shared fence device

In this example I've defined a WTI power switch. I defined:

- the name of the switch in the cluster
- the ip address
- login
- password
- power delay



Create a Fence Device ✕-- Select a fence device -- ⬆

Fence type

WTI Power Switch

Name

summit-wti-switch

IP address or hostname

172.31.100.254

Login

root

Password

●●●●●●●●

Password Script (optional)

Force command prompt

Power wait (seconds)

5

Submit


Cancel

QDisk

Configure

Hostname

ster is unknown.



Name	Fence Type	Nodes Using	Hostname
☐ summit-wti-switch	WTI Power Device	0	172.31.100.254

Select an item to view details

Set up per-node fence devices

Now we need to assign a fence device to each node.

Click on Nodes tab

Click on a host

Add a fence device

Assign a power port (in this case)

Repeat for other node

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Cluster Management - Mozilla Firefox

FileEditViewHistoryBookmarksToolsHelp

←→↺⌂🔍

redhat.comhttps://neuromancer.tc.redhat.com:8084/cluster/summit-cluster/molly

Google

Most VisitedRed HatRed Hat MagazineRed Hat NetworkRed Hat Support

Cluster ManagementTest Page for the Apache H...

Homebase

Manage Clusters

summit-cluster

NodesServicesResourcesFailover DomainsFence DevicesQDiskConfigure

+ Add↺ Reboot↻ Join Cluster↻ Leave Cluster⌫ Delete

!	Node Name	Node ID	Votes	Status	Hostname
☐	molly.tc.redhat.com	2	1	Cluster Member	molly.tc.redhat.com
☐	case.tc.redhat.com	1	1	Cluster Member	case.tc.redhat.com

molly.tc.redhat.com

Status Cluster Member

⌫↺↻↻

Services

Failover Domains

Fence Devices

PrimarySecondary

Fence type

WTI Power Switch

Name

summit-wti-switch


IP address or hostname

172.31.100.254

Login

root

Done



Password Script (optional)

Force command prompt

Power wait (seconds)

0

Port

1

SSHUse
SSH☐Path to
ssh
identity
file

Remove this instance

Add instance

Remove this fence device

-- Select a fence device --

Add a fence to this level

Apply

-- Select a fence device --

Add a fence to this level

Set up failover domains

Define a failover domain to set the priority of a host for a given service.

In this example, we'll define a failover domain which makes it more likely that the service we define later will come up on case.


Go to the failover domain tab and click “Add.”

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





Set up failover domains

I've set up the failover domain name as “prefer_case”

I've set the failover domain to be prioritized

I have not set the failover domain to be restricted

I have not set the “No failback” - I want the service to run on case whenever it's available

I've checked the membership boxes next to molly and case, and set case to have a lower value than molly

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Cluster Management - Mozilla Firefox

FileEditViewHistoryBookmarksToolsHelp

←→⌵↺⌵🏠🔍

redhat.comhttps://neuromancer.tc.redhat.com:8084/cluster/summit-cluster/failov

Google

Most Visited

Red HatRed Hat MagazineRed Hat NetworkRed Hat Support

Cluster Management

+

cluster
management

AboutLogout

Homebase

Clusters

summit-cluster

✓Deleting failover domain

Homebase

Manage Clusters

summit-cluster

QDisk

Configure

Restricted

Cluster is unknown.

Create a Failover Domain

Nameprefer_case

☒PrioritizedOrder the nodes to which services failover.

☐RestrictedService can run only on nodes specified.


☐No FailbackDo not send service back to 1st priority node when it becomes available again.

	Member	Priority
	molly.tc.redhat.com	☒ 10
	case.tc.redhat.com	☒ 1

Create

Cancel

Done



Define Resources for the Service

You'll typically define three resources for the clustered web application.

- Virtual ip address
- Start/stop script
- Shared storage (not necessary since it's shared)




Define Resources for the Service

Define a virtual ip address.

- Resources
 - IP address
 - Monitor link
 - Sleep time





Cluster Management - Mozilla Firefox

FileEditViewHistoryBookmarksToolsHelp

←→↺⌂

redhat.com

https://neuromancer.tc.redhat.com:8084/cluster/summit-cluster/resol

☆

Google

Most Visited

Red Hat

Red Hat Magazine

Red Hat Network

Red Hat Support

Cluster Management

Test Page for the Apache H...

cluster

management

AboutLogout

Homebase

Clusters

summit-cluster

Homebase

Manage Clusters

summit-cluster

Nodes

Services

Resources

Failover Domains

Fence Devices

QDisk

Configure

+ Add

× Delete

	Name/IP	Type	In Use
☐	172.31.100.253	IP Address	

Select an item to view details

Done

Define Resources for the Service

Define a script resource for httpd

- Note that the httpd service should **not** be chkconfig'd on on the nodes! We want the clustering software to start it, not SysV init.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Define Resources for the Service

Define a script resource


- Resources
 - Resource name
 - Server root
 - Config file
 - Apache options
 - Shutdown delay


SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT







	Name/IP	Type	In Use
☐	172.31.100.253	IP Address	
☐	summit-httpd-script	Script	

[Select an item to view details](#)

Define Resources for the Service

Shared storage could be added here as well, but since we mounted it on each host via `/etc/fstab`, there's no need to define it in the cluster.



Define a Service

We're going to define a web server service to start and listen to 172.31.100.253. We need to edit the `/etc/httpd/conf/httpd.conf` file on each server so that the web server listens on that ip address.



```
root@molly:~
File Edit View Search Terminal Help
MaxRequestsPerChild 0
</IfModule>

#
# Listen: Allows you to bind Apache to specific IP addresses and/or
# ports, in addition to the default. See also the <VirtualHost>
# directive.
#
# Change this to Listen on specific IP addresses as shown below to
# prevent Apache from glomming onto all bound IP addresses (0.0.0.0)
#
#Listen 12.34.56.78:80
Listen 172.31.100.253:80

#
# Dynamic Shared Object (DSO) Support
#
# To be able to use the functionality of a module which was built as a DSO you
# have to place corresponding 'LoadModule' lines at this location so the
# directives contained in it are actually available _before_ they are used.
# Statically compiled modules (those listed by 'httpd -l') do not need
# to be loaded here.
#
"/etc/httpd/conf/httpd.conf" 1008L, 34417C written
```

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Define a Service

Now we go to the Services tab and define a web server service.

Add the resources to the service

Choose:

- Auto-start
- Failover domain
- Recovery policy



Cluster Management - Mozilla Firefox

FileEditViewHistoryBookmarksToolsHelp

←→⌵↺⌵🏠🔍redhat.comhttps://neuromancer.tc.redhat.com:8084/cluster/summit-cluster/servi★

Google

Most Visited

Red Hat

Red Hat Magazine

Red Hat Network

Red Hat Support

Cluster Management

+

▼

cluster

management

AboutLogout

Homebase

Clusters

summit-cluster

Homebase

Manage Clusters

● summit-cluster

Nodes

Services

Resources

Failover Domains

Fence Devices

QDisk

Configure

⊕ Add

▶ Start


↺ Restart


■ Disable


✕ Delete

!	Name	Status	Autostart	Failover Domain
No nodes from this cluster could be contacted. The status of this cluster is unknown.				

https://neuromancer.tc.redhat.com:8084/cluster/summit-cluster/services#







Define a Service

You can also check the logs on each host

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



```
root@case:~  
File Edit View Search Terminal Help  
Jun 19 23:41:31 case modcluster: Updating cluster.conf  
Jun 19 23:41:31 case modcluster: Updating cluster version  
Jun 19 23:41:32 case rgmanager[2232]: Reconfiguring  
Jun 19 23:41:32 case rgmanager[2232]: Loading Service Data  
Jun 19 23:41:32 case rgmanager[2232]: Stopping changed resources.  
Jun 19 23:41:32 case rgmanager[2232]: Restarting changed resources.  
Jun 19 23:41:32 case rgmanager[2232]: Starting changed resources.  
Jun 19 23:44:05 case modcluster: Updating cluster.conf  
Jun 19 23:44:05 case modcluster: Updating cluster version  
Jun 19 23:44:06 case rgmanager[2232]: Reconfiguring  
Jun 19 23:44:06 case rgmanager[2232]: Loading Service Data  
Jun 19 23:44:07 case rgmanager[2232]: Stopping changed resources.  
Jun 19 23:44:07 case rgmanager[2232]: Restarting changed resources.  
Jun 19 23:44:07 case rgmanager[2232]: Starting changed resources.  
Jun 19 23:44:07 case rgmanager[2232]: Initializing service:summit-web-service  
Jun 19 23:44:07 case bash[12650]: Executing /etc/rc.d/init.d/httpd stop  
Jun 19 23:44:07 case rgmanager[2232]: Starting stopped service service:summit-we  
b-service  
Jun 19 23:44:07 case bash[12779]: Adding IPv4 address 172.31.100.253/24 to eth0  
Jun 19 23:44:09 case avahi-daemon[1457]: Registering new address record for 172.  
31.100.253 on eth0.IPv4.  
Jun 19 23:44:11 case bash[12868]: Executing /etc/rc.d/init.d/httpd start  
Jun 19 23:44:11 case rgmanager[2232]: Service service:summit-web-service started  
  
Jun 19 23:42:35 molly modcluster: Updating cluster.conf  
Jun 19 23:42:35 molly modcluster: Updating cluster version  
Jun 19 23:42:35 molly rgmanager[2166]: Reconfiguring  
Jun 19 23:42:35 molly rgmanager[2166]: Loading Service Data  
Jun 19 23:42:36 molly rgmanager[2166]: Stopping changed resources.  
Jun 19 23:42:36 molly rgmanager[2166]: Restarting changed resources.  
Jun 19 23:42:36 molly rgmanager[2166]: Starting changed resources.  
Jun 19 23:42:36 molly rgmanager[2166]: Initializing service:summit-web-service  
Jun 19 23:42:37 molly bash[11598]: Executing /etc/rc.d/init.d/httpd stop
```

Manage the services

You can migrate services from one host to another via the Conga web UI or from the command line

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Manage the services


In the Conga web UI, click on the service, click the dropdown and choose which host to start the service on:


SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT







Manage the services

You can also look in the log files of each host to check the progress of the migration:

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



```
root@case:~  
File Edit View Search Terminal Help  
  
Jun 20 01:32:57 case modcluster: Starting service: summit-web-service on node  
Jun 20 01:32:58 case rgmanager[1887]: Stopping service service:summit-web-service  
Jun 20 01:32:58 case bash[4360]: Executing /etc/rc.d/init.d/httpd stop  
Jun 20 01:32:59 case bash[4434]: Removing IPv4 address 172.31.100.253/24 from eth0  
Jun 20 01:32:59 case avahi-daemon[1601]: Withdrawing address record for 172.31.100.253 on eth0.  
Jun 20 01:32:59 case rgmanager[1887]: Service service:summit-web-service is stopped  
Jun 20 01:33:03 case rgmanager[1887]: Service service:summit-web-service is now running on member 2
```

```
Jun 20 01:31:29 molly rgmanager[1810]: Starting stopped service service:summit-web-service  
Jun 20 01:31:29 molly bash[4608]: Adding IPv4 address 172.31.100.253/24 to eth0  
Jun 20 01:31:31 molly avahi-daemon[1533]: Registering new address record for 172.31.100.253 on eth0.IPv4.  
Jun 20 01:31:33 molly bash[4692]: Executing /etc/rc.d/init.d/httpd start  
Jun 20 01:31:33 molly rgmanager[1810]: Service service:summit-web-service started
```

Manage the services

You can also migrate the service from the command line:

- clustat
- clusvcadm



```

root@case:~
File Edit View Search Terminal Help
[root@case ~]# clustat
Cluster Status for summit-cluster @ Sun Jun 20 01:34:50 2010
Member Status: Quorate

Member Name          ID   Status
-----
case.tc.redhat.com    1   Online, Local, rgmanager
molly.tc.redhat.com   2   Online, rgmanager

Service Name          Owner (Last)          State
-----
service:summit-web-service molly.tc.redhat.com started
[root@case ~]# clusvcadm -r service:summit-web-service -m case.tc.redhat.com
Trying to relocate service:summit-web-service to case.tc.redhat.com...Success
service:summit-web-service is now running on case.tc.redhat.com
[root@case ~]# clustat
Cluster Status for summit-cluster @ Sun Jun 20 01:35:52 2010
Member Status: Quorate

Member Name          ID   Status
-----
case.tc.redhat.com    1   Online, Local, rgmanager
molly.tc.redhat.com   2   Online, rgmanager

Service Name          Owner (Last)          State
-----
service:summit-web-service case.tc.redhat.com started
[root@case ~]#

```

```

~
...
vice service:summit-web-service is now
...
rting stopped service service:summit-w
...
Pv4 address 172.31.100.253/24 to eth0
Registering new address record for 172
...
g /etc/rc.d/init.d/httpd start
vice service:summit-web-service starte
...
g /etc/rc.d/init.d/httpd status
g /etc/rc.d/init.d/httpd status
g /etc/rc.d/init.d/httpd status
opping service service:summit-web-servi
...
g /etc/rc.d/init.d/httpd stop
IPv4 address 172.31.100.253/24 from e
...
Jun 20 01:33:41 molly avahi-daemon[1533]: Withdrawing address record for 172.31.
100.253 on eth0.
Jun 20 01:33:41 molly rgmanager[1810]: Service service:summit-web-service is sto
pped

```

More Information

Red Hat Online Documentation:

- <http://www.redhat.com/docs/>

Red Hat Enterprise Clustering and Storage Management:

- <http://bit.ly/9KEDhZ>

Red Hat's Cluster Wiki:

- <http://sources.redhat.com/cluster/wiki/>

Red Hat Mailing Lists, e.g.

- <https://listman.redhat.com/mailman/listinfo/linux-cluster>

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Thank You!

If you liked today's presentation, please let us know!

You can follow up with us:

Thomas - thomas@redhat.com, choirboy on #rhel on Freenode, thomasdcameron on Twitter

Lon - lh@redhat.com

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT





SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

