

RED HAT :: CHICAGO :: 2009

SUMMIT

FOLLOW US:

[TWITTER.COM/REDHATSUMMIT](https://twitter.com/redhatsummit)

TWEET ABOUT US:

ADD #SUMMIT AND/OR #JBOSSWORLD TO THE END
OF YOUR EVENT-RELATED TWEET

presented by



Clustering Virtual Machines and Applications with Red Hat Enterprise Linux

Thomas Cameron, RHCE, RHCX, CNE, MCSE, MCT
Solutions Architect Team Lead, Red Hat
Lon Hohberger

Senior Cluster Software Engineer, Red Hat

September 2, 2009

This Presentation

The updated version of this presentation is available at
<http://people.redhat.com/tcameron>

If you like this presentation, please take the survey at
<http://www.redhat.com/summit-survey> and tell your friends!

Agenda

Red Hat and Virtualization

Architecture

Configure Raw Storage and iSCSI

- On the Target

- On the Client

Define dom0 Cluster

- Conga

Configure Shared Storage for dom0 Machines

- Conga

Install domU machines

Agenda

Set up cluster of domU machines

Set up Application Cluster on domU machines

Questions

Red Hat and Virtualization

Red Hat leads the way in Open Source virtualization

Highly successful entry into virtualization with Xen at the release of RHEL 5.0

Acquired Qumranet, the company which developed KVM
September 2008

KVM included in RHEL 5.4 – more choice, more flexibility!

Announcing Red Hat Enterprise Virtualization (RHEV), a KVM based hypervisor and enterprise virtualization management system at Summit this year.

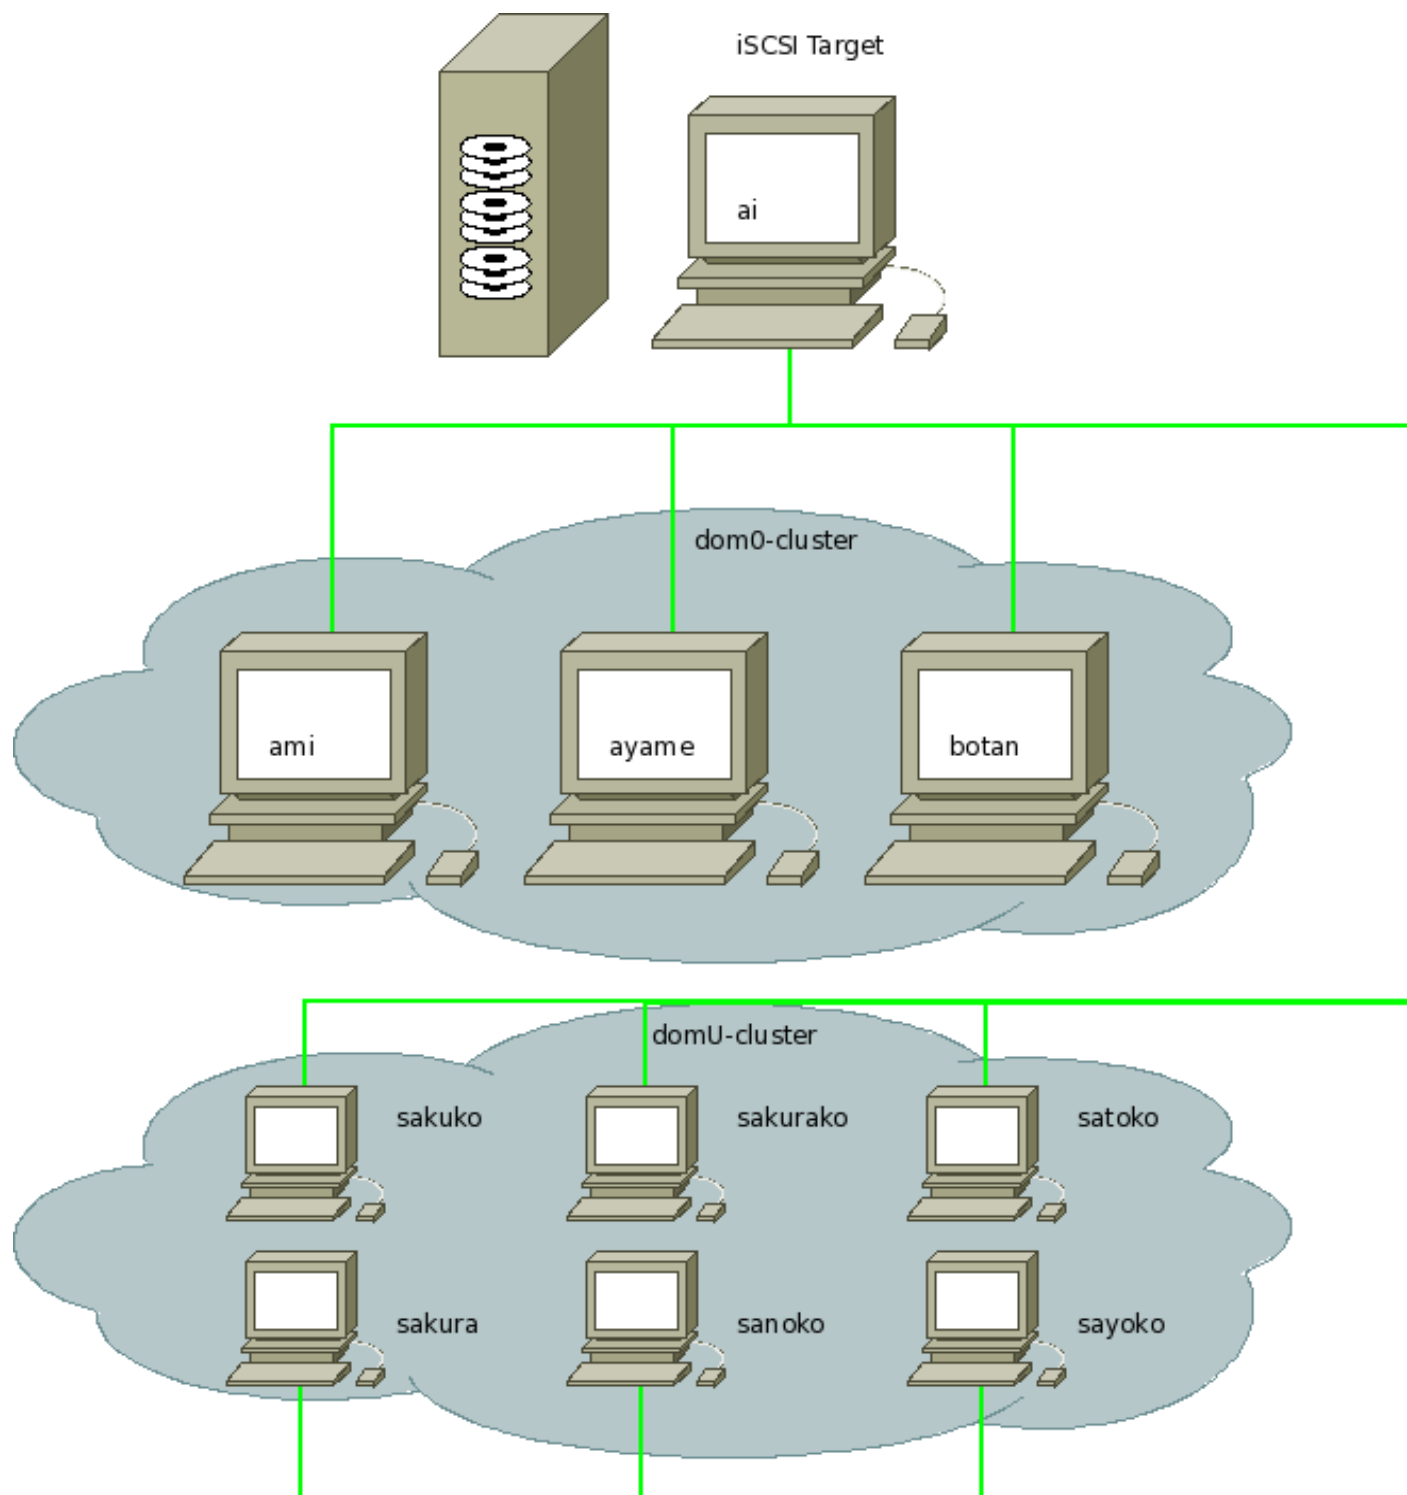
Architecture

ai - iSCSI target and Conga server. Not part of the cluster, just storage and management. Kickstarted with cluster-storage and clustering groups and scsi-target-utils package

ami, ayame and botan - iSCSI initiators, dom0 cluster members. Kickstarted with virtualization, cluster-storage and clustering groups and iscsi-initiator-utils package

sakuko, sakura, sakurako, sanako, satoko and sayoko - domU machines. Kickstarted with cluster-storage and clustering groups. Running an Apache cluster.

All systems are on a gigabit ethernet network



Configure storage and iSCSI on ai

Create the physical volume, volume group and logical volume. In this case, /dev/sdb1 is a newly created, blank partition of 1TiB:

```
pvcreate /dev/sdb1
```

```
vgcreate VolGroup01 /dev/sdb1
```

```
lvcreate -n VMStorage -L 100GB VolGroup01
```

Configure storage and iSCSI on ai

Make sure that the iSCSI target software is installed

```
yum -y groupinstall clustering cluster-storage
```

```
yum -y install scsi-target-utils
```

```
chkconfig tgtd on
```


Configure storage and iSCSI on ai

Configure the target

Configure the `/etc/tgtd/target.conf` to look something like this:

```
<target iqn.2009-06.com.redhat.dfw.salab:disk1>  
    backing-store /dev/mapper/VolGroup01-VMStorage  
    initiator-address 10.15.183.6  
    initiator-address 10.15.183.7  
    initiator-address 10.15.183.8  
</target>
```

Configure storage and iSCSI on ai

Stop and start the tgtd service on the iSCSI target

```
service tgtd stop
```

```
sleep 5
```

```
service tgtd start
```

Configure storage and iSCSI on ai

Check that your target is available

```
tgtadm --lld iscsi --op show --mode target
```



```
root@ai:~  
File Edit View Terminal Help  
[root@ai ~]# tgtadm --lld iscsi --op show --mode target  
Target 1: iqn.2009-06.com.redhat.dfw.salab:disk1  
  System information:  
    Driver: iscsi  
    State: ready  
  I_T nexus information:  
    I_T nexus: 26  
      Initiator: iqn.1994-05.com.redhat:88647efd3573  
      Connection: 0  
      IP Address: 10.15.183.6  
    I_T nexus: 27  
      Initiator: iqn.1994-05.com.redhat:d6eff839a065  
      Connection: 0  
      IP Address: 10.15.183.8  
    I_T nexus: 28  
      Initiator: iqn.1994-05.com.redhat:437eb150453a  
      Connection: 0  
      IP Address: 10.15.183.7  
  LUN information:  
    LUN: 0  
      Type: controller  
      SCSI ID: deadbeaf1:0  
      SCSI SN: beaf10  
      Size: 0 MB  
      Online: Yes  
      Removable media: No  
      Backing store: No backing store  
    LUN: 1  
      Type: disk  
      SCSI ID: deadbeaf1:1  
      SCSI SN: beaf11  
      Size: 107374 MB  
      Online: Yes  
      Removable media: No  
      Backing store: /dev/mapper/VolGroup01-VMStorage  
  Account information:  
  ACL information:  
    10.15.183.6  
    10.15.183.7  
    10.15.183.8  
[root@ai ~]#
```

Configure storage and iSCSI on initiators

On each of the dom0 servers (ami, ayame and botan)

```
yum -y groupinstall clustering cluster-storage
```

```
yum -y install iscsi-initiator-utils
```

```
chkconfig iscsi on
```

```
chkconfig iscsid on
```

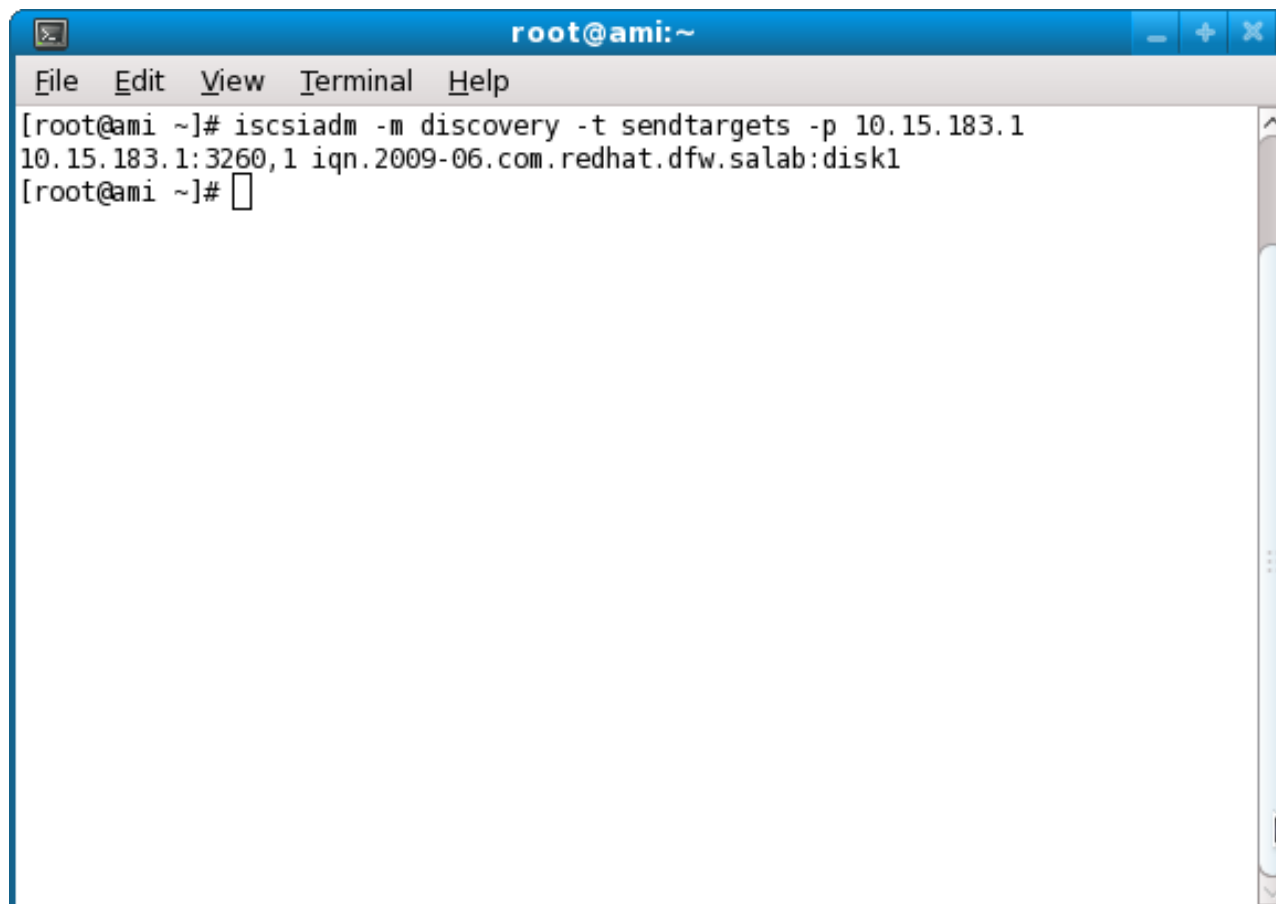
```
service iscsid restart
```

```
service iscsi restart
```

Configure storage and iSCSI on initiators

Discover the iSCSI target

```
iscsiadm -m discovery -t sendtargets -p 10.15.183.1
```

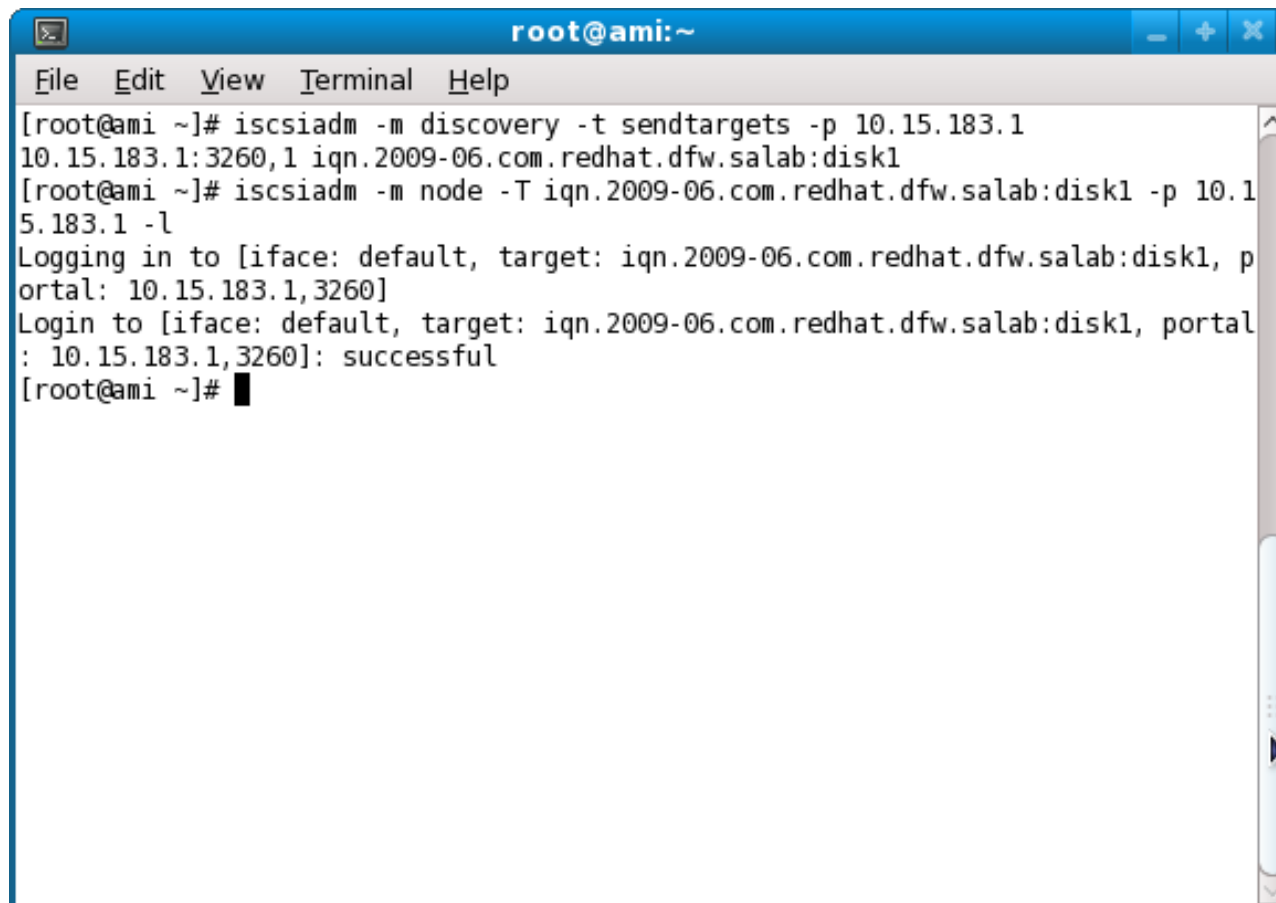
A terminal window titled 'root@ami:~' with a menu bar containing 'File', 'Edit', 'View', 'Terminal', and 'Help'. The terminal shows the command '[root@ami ~]# iscsiadm -m discovery -t sendtargets -p 10.15.183.1' and its output '10.15.183.1:3260,1 iqn.2009-06.com.redhat.dfw.salab:disk1'. The prompt '[root@ami ~]#' is followed by a cursor.

```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# iscsiadm -m discovery -t sendtargets -p 10.15.183.1  
10.15.183.1:3260,1 iqn.2009-06.com.redhat.dfw.salab:disk1  
[root@ami ~]#
```

Configure storage and iSCSI on initiators

Log in to the target (command below is actually all one command, it is on multiple lines for ease of visibility)

```
iscsiadm -m node \  
-T iqn.2009-06.com.redhat.dfw.salab:disk1 \  
-p 10.15.183.1 -l
```



```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# iscsiadm -m discovery -t sendtargets -p 10.15.183.1  
10.15.183.1:3260,1 iqn.2009-06.com.redhat.dfw.salab:disk1  
[root@ami ~]# iscsiadm -m node -T iqn.2009-06.com.redhat.dfw.salab:disk1 -p 10.15.183.1 -l  
Logging in to [iface: default, target: iqn.2009-06.com.redhat.dfw.salab:disk1, portal: 10.15.183.1,3260]  
Login to [iface: default, target: iqn.2009-06.com.redhat.dfw.salab:disk1, portal: 10.15.183.1,3260]: successful  
[root@ami ~]#
```


Configure storage and iSCSI on initiators

Run “fdisk -l” or “cat /proc/partitions” to verify that your system “sees” the new disk

There should be no partition table on e.g. /dev/sdb

```
root@ami:~  
File Edit View Terminal Help  
10.15.183.1:3260,1 iqn.2009-06.com.redhat.dfw.salab:disk1  
[root@ami ~]# iscsiadm -m node -T iqn.2009-06.com.redhat.dfw.salab:disk1 -p 10.15.183.1 -l  
Logging in to [iface: default, target: iqn.2009-06.com.redhat.dfw.salab:disk1, portal: 10.15.183.1,3260]  
Login to [iface: default, target: iqn.2009-06.com.redhat.dfw.salab:disk1, portal: 10.15.183.1,3260]: successful  
[root@ami ~]# fdisk -l  
  
Disk /dev/sda: 1000.2 GB, 1000204886016 bytes  
255 heads, 63 sectors/track, 121601 cylinders  
Units = cylinders of 16065 * 512 = 8225280 bytes  


| Device    | Boot | Start | End    | Blocks    | Id | System               |
|-----------|------|-------|--------|-----------|----|----------------------|
| /dev/sda1 | *    | 1     | 13     | 104391    | 83 | Linux                |
| /dev/sda2 |      | 14    | 281    | 2152710   | 82 | Linux swap / Solaris |
| /dev/sda3 |      | 282   | 121601 | 974502900 | 8e | Linux LVM            |

  
Disk /dev/sdb: 107.3 GB, 107374182400 bytes  
255 heads, 63 sectors/track, 13054 cylinders  
Units = cylinders of 16065 * 512 = 8225280 bytes  
  
Disk /dev/sdb doesn't contain a valid partition table  
[root@ami ~]#
```

Configure storage and iSCSI on initiators

Create a partition on the LUN using fdisk

```
fdisk /dev/sdb
```

Choose n for new

Choose p for primary

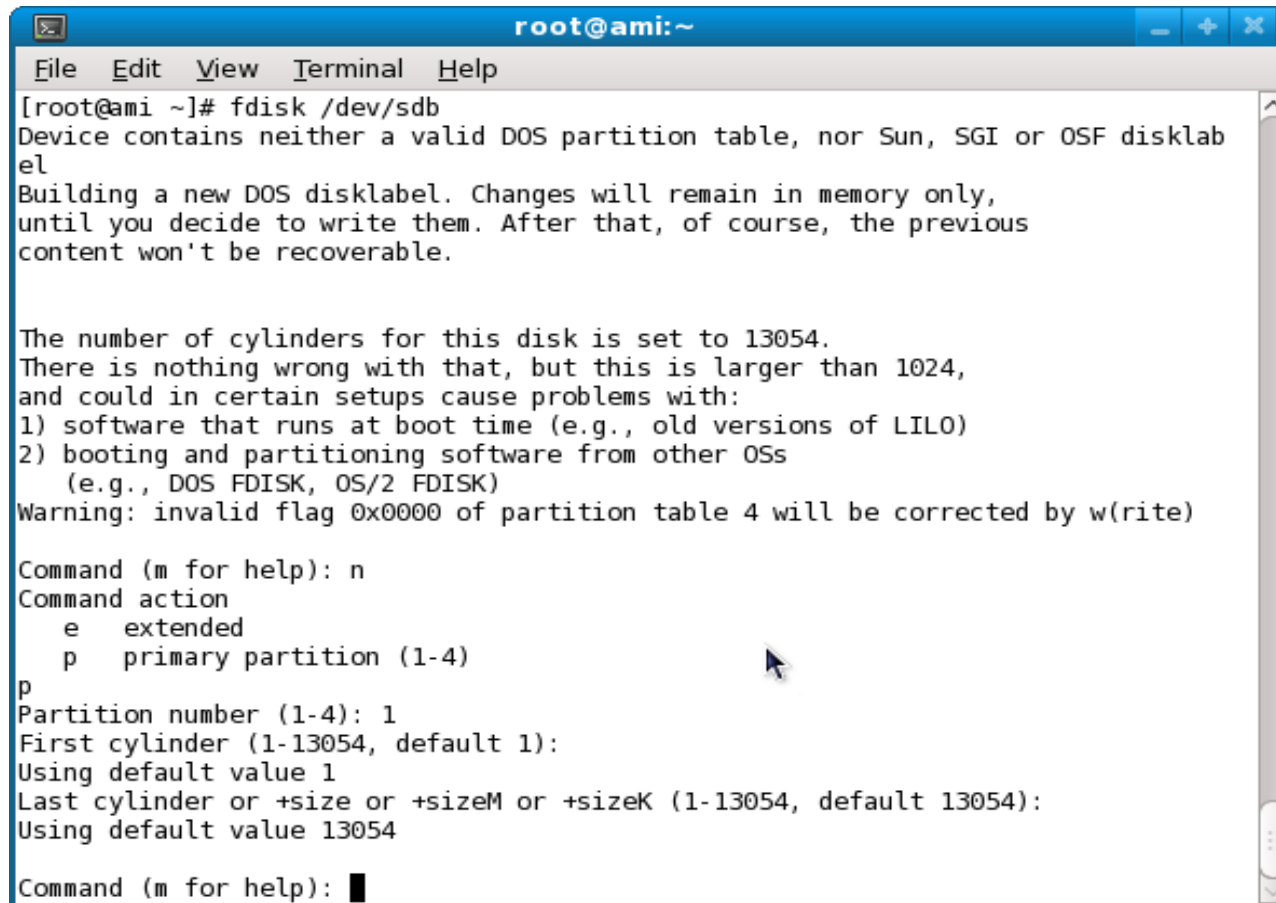
Choose 1 for first partition

Hit [enter] to start at the first cylinder

Hit [enter] to end at the last cylinder

Hit w to write changes to disk

As root, run “partprobe” and “fdisk -l” on the other nodes and verify that the new partition shows up.



```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# fdisk /dev/sdb  
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel  
Building a new DOS disklabel. Changes will remain in memory only,  
until you decide to write them. After that, of course, the previous  
content won't be recoverable.  
  
The number of cylinders for this disk is set to 13054.  
There is nothing wrong with that, but this is larger than 1024,  
and could in certain setups cause problems with:  
1) software that runs at boot time (e.g., old versions of LILO)  
2) booting and partitioning software from other OSs  
   (e.g., DOS FDISK, OS/2 FDISK)  
Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)  
  
Command (m for help): n  
Command action  
   e   extended  
   p   primary partition (1-4)  
p  
Partition number (1-4): 1  
First cylinder (1-13054, default 1):  
Using default value 1  
Last cylinder or +size or +sizeM or +sizeK (1-13054, default 13054):  
Using default value 13054  
Command (m for help):
```

Configure storage and iSCSI on initiators

Make sure that the ricci service is installed and running on all the dom0 nodes.

```
service ricci status
```

Configure the Cluster Using Conga

On the management station (ai), make sure that the luci package is installed

```
yum -y install luci
```

```
chkconfig luci on
```

Don't start the luci service yet!

Add an administrator password

```
luci_admin init
```

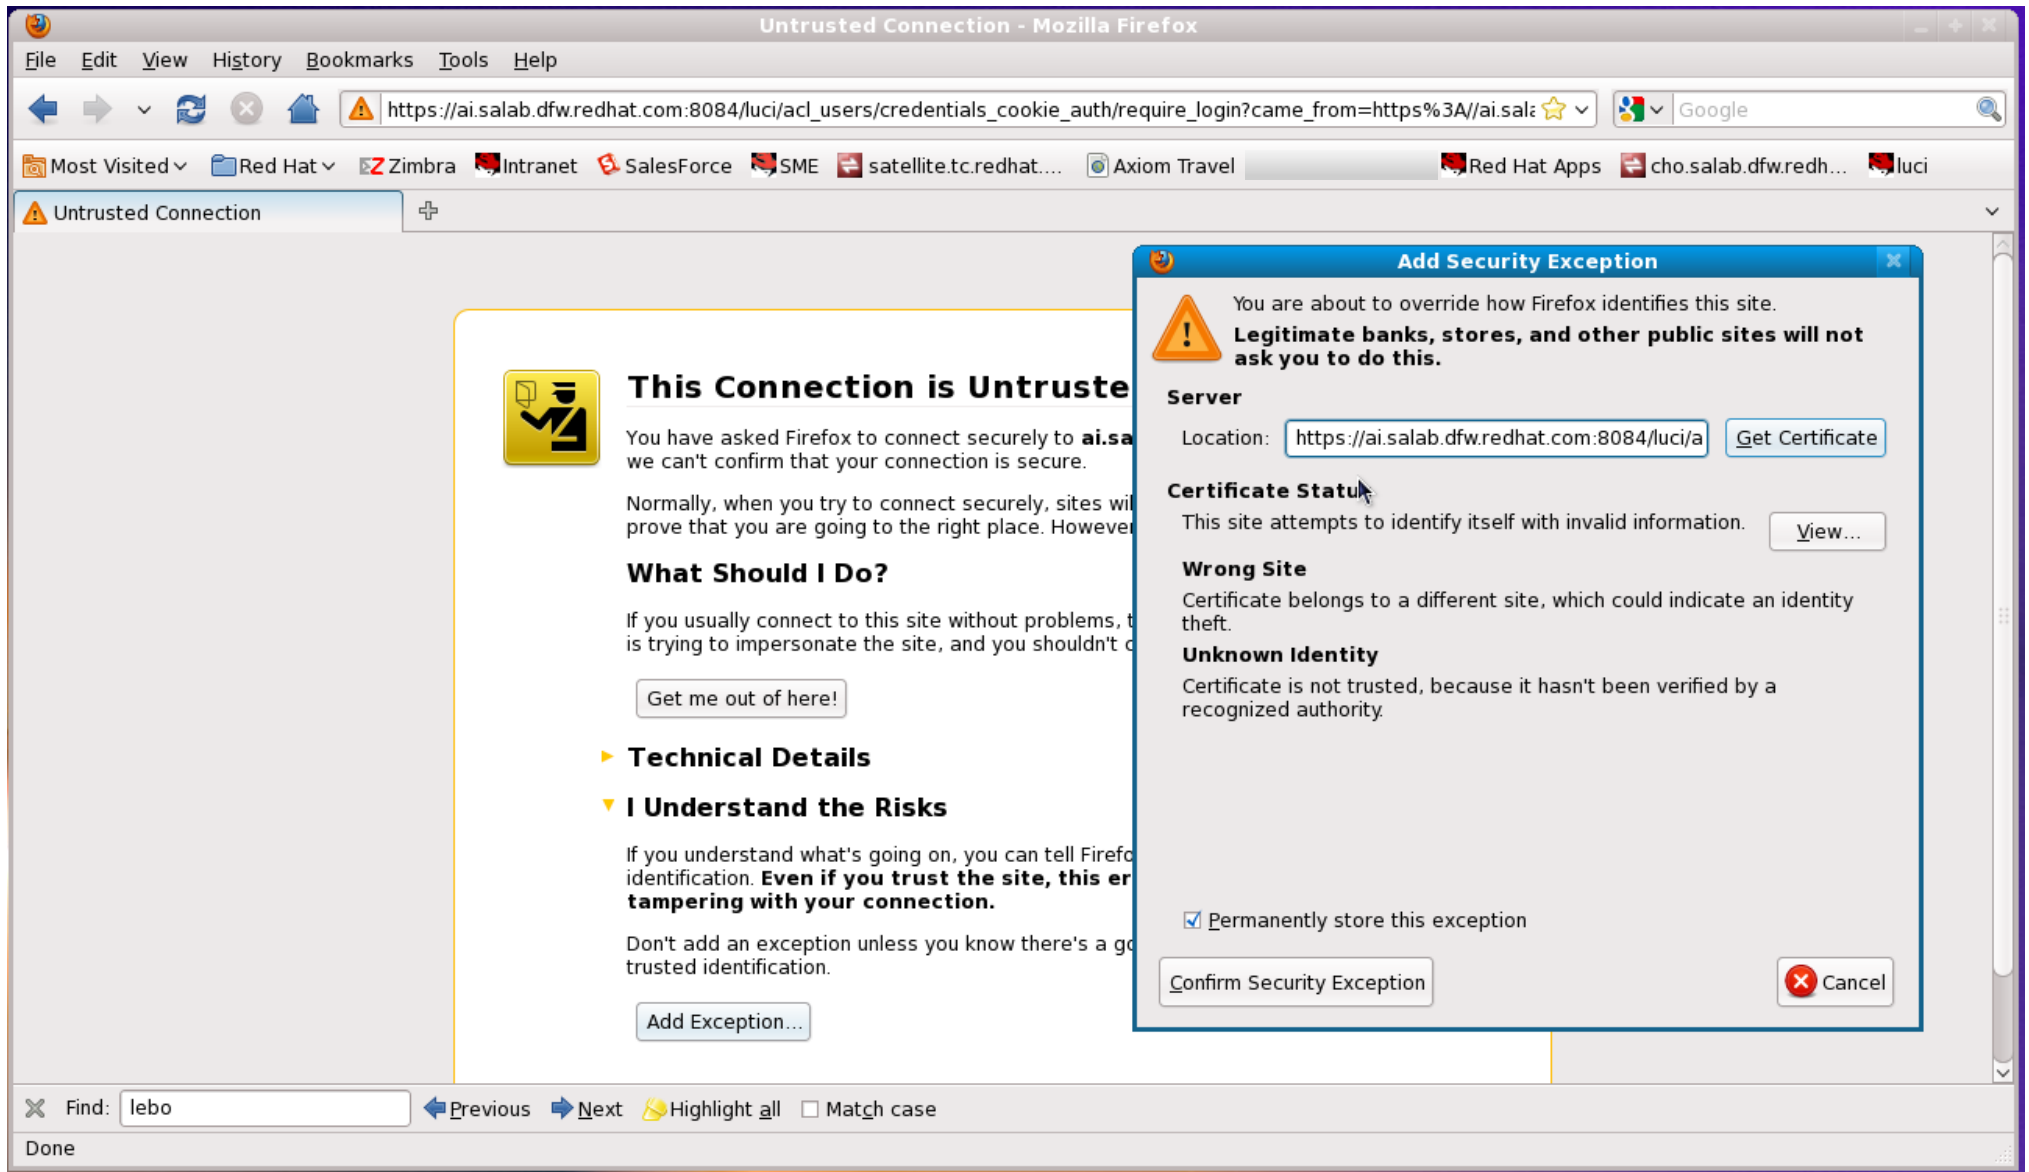
```
service luci start
```

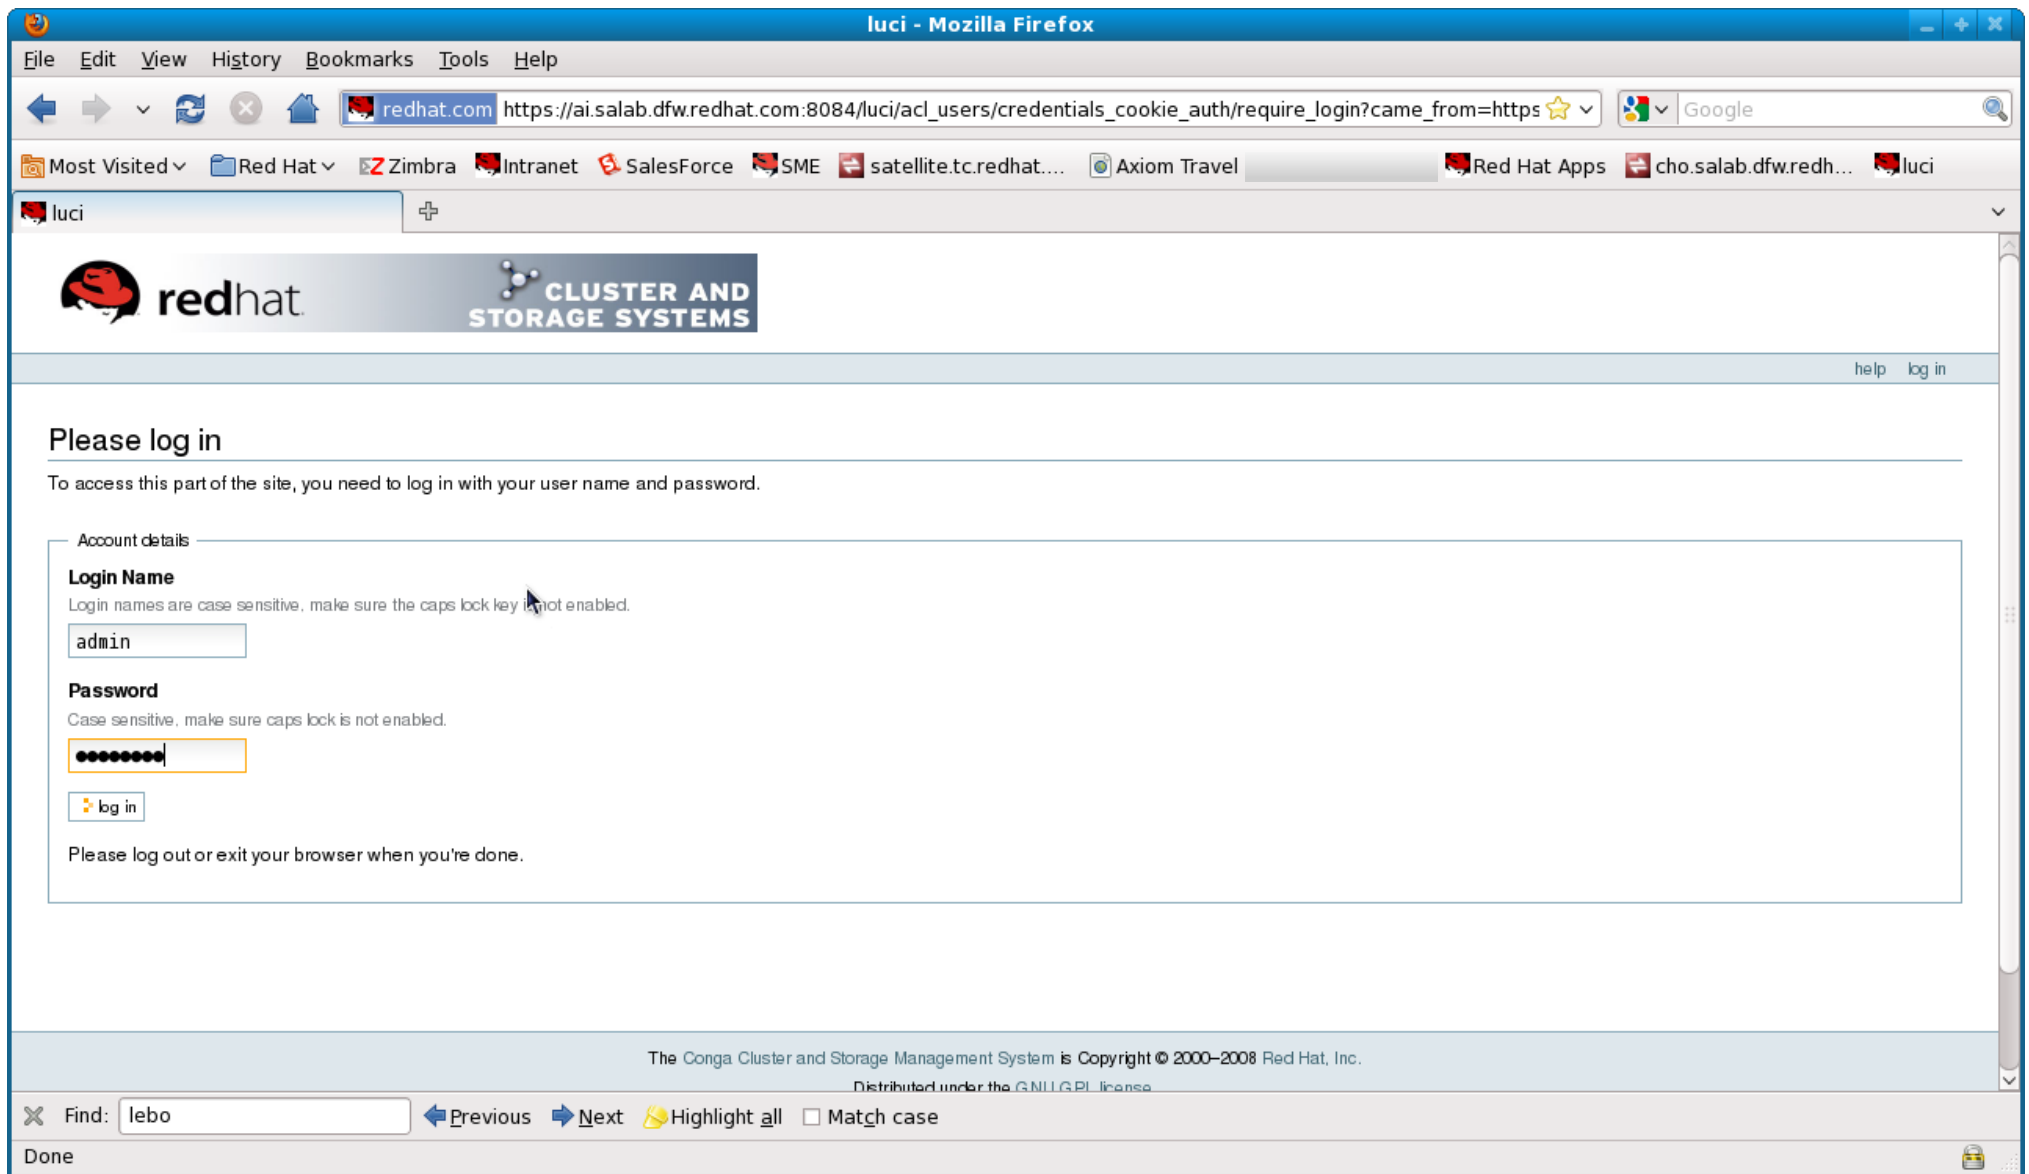
```
root@ai:~  
File Edit View Terminal Help  
[root@ai ~]# chkconfig luci on  
[root@ai ~]# service luci status  
luci is stopped  
[root@ai ~]# luci_admin init  
Initializing the luci server  
  
Creating the 'admin' user  
  
Enter password:  
Confirm password:  
  
Please wait...  
The admin password has been successfully set.  
Generating SSL certificates...  
The luci server has been successfully initialized  
  
You must restart the luci server for changes to take effect.  
Run "service luci restart" to do so  
  
[root@ai ~]# service luci restart  
Shutting down luci: [ OK ]  
Starting luci: Generating https SSL certificates... done [ OK ]  
  
Point your web browser to https://ai.salab.dfw.redhat.com:8084 to access luci  
[root@ai ~]#
```

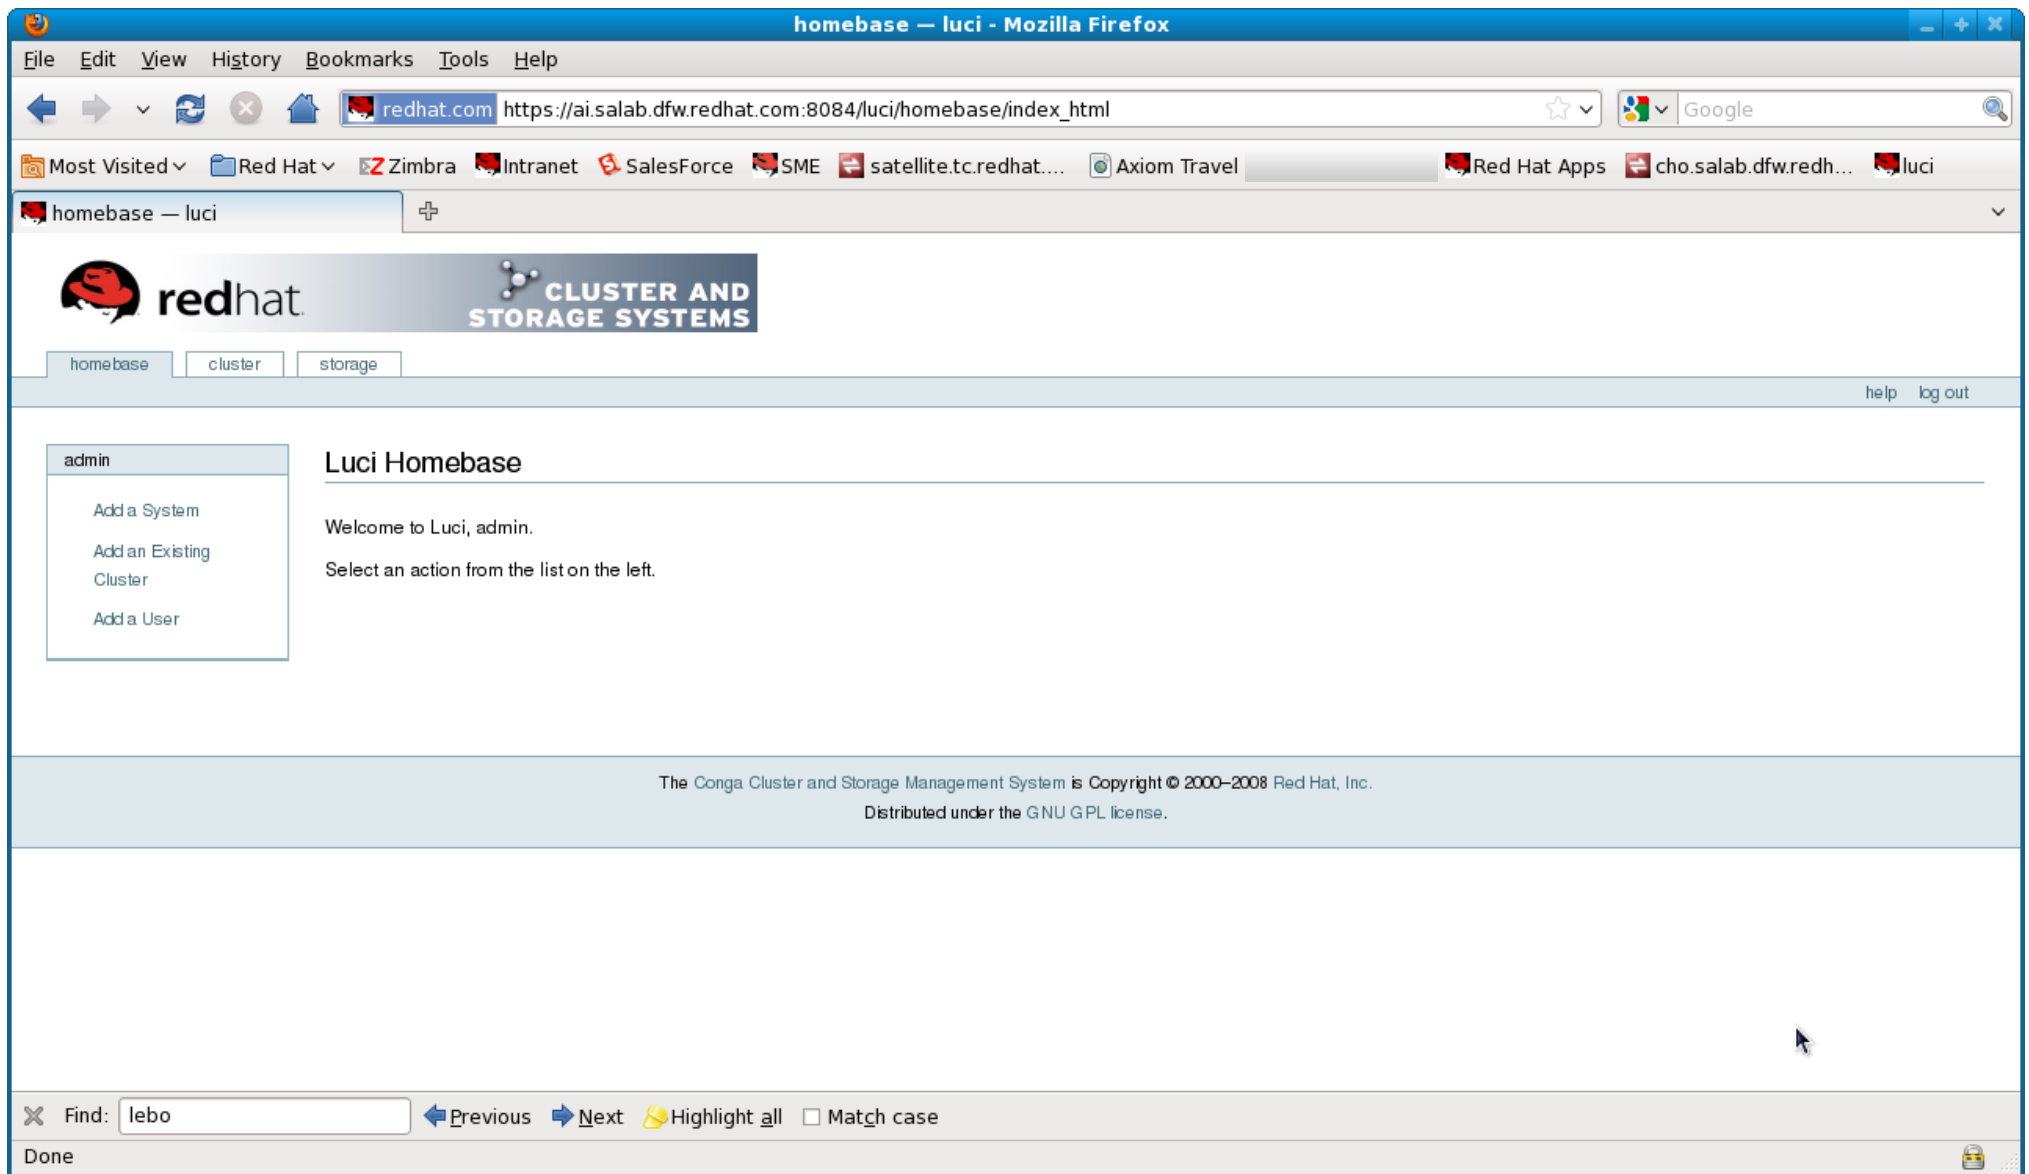
Configure the Cluster Using Conga

Open the URL listed by the “service luci start” command in your web browser

You will be warned about the self-signed certificate, it's fine to accept it assuming you trust your network







Configure the Cluster Using Conga

Choose the “Cluster” tab and create a new cluster

Cluster name

Node names and passwords

Download or use locally installed

Enable shared storage since we'll be using GFS

Reboot nodes

Check the box if the passwords are all the same

Luci — cluster — Deploy a cluster - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=6

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — Deploy a cluster

redhat CLUSTER AND STORAGE SYSTEMS







homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

Create a new cluster

Cluster Name

| Node Hostname | Root Password | Key ID |
|---|--|---|
| <input type="text" value="ami.salab.dfw.redhat.com"/> | <input type="password" value="....."/> |   |
| <input type="text" value="ayame.salab.dfw.redhat.com"/> | <input type="password" value="....."/> |   |
| <input type="text" value="botan.salab.dfw.redhat.com"/> | <input type="password" value="....."/> |   |

☒ Download packages
☐ Use locally installed packages.

☒ Enable Shared Storage Support
☒ Reboot nodes before joining cluster
☒ Check if node passwords are identical.

Find: Previous Next Highlight all Match case

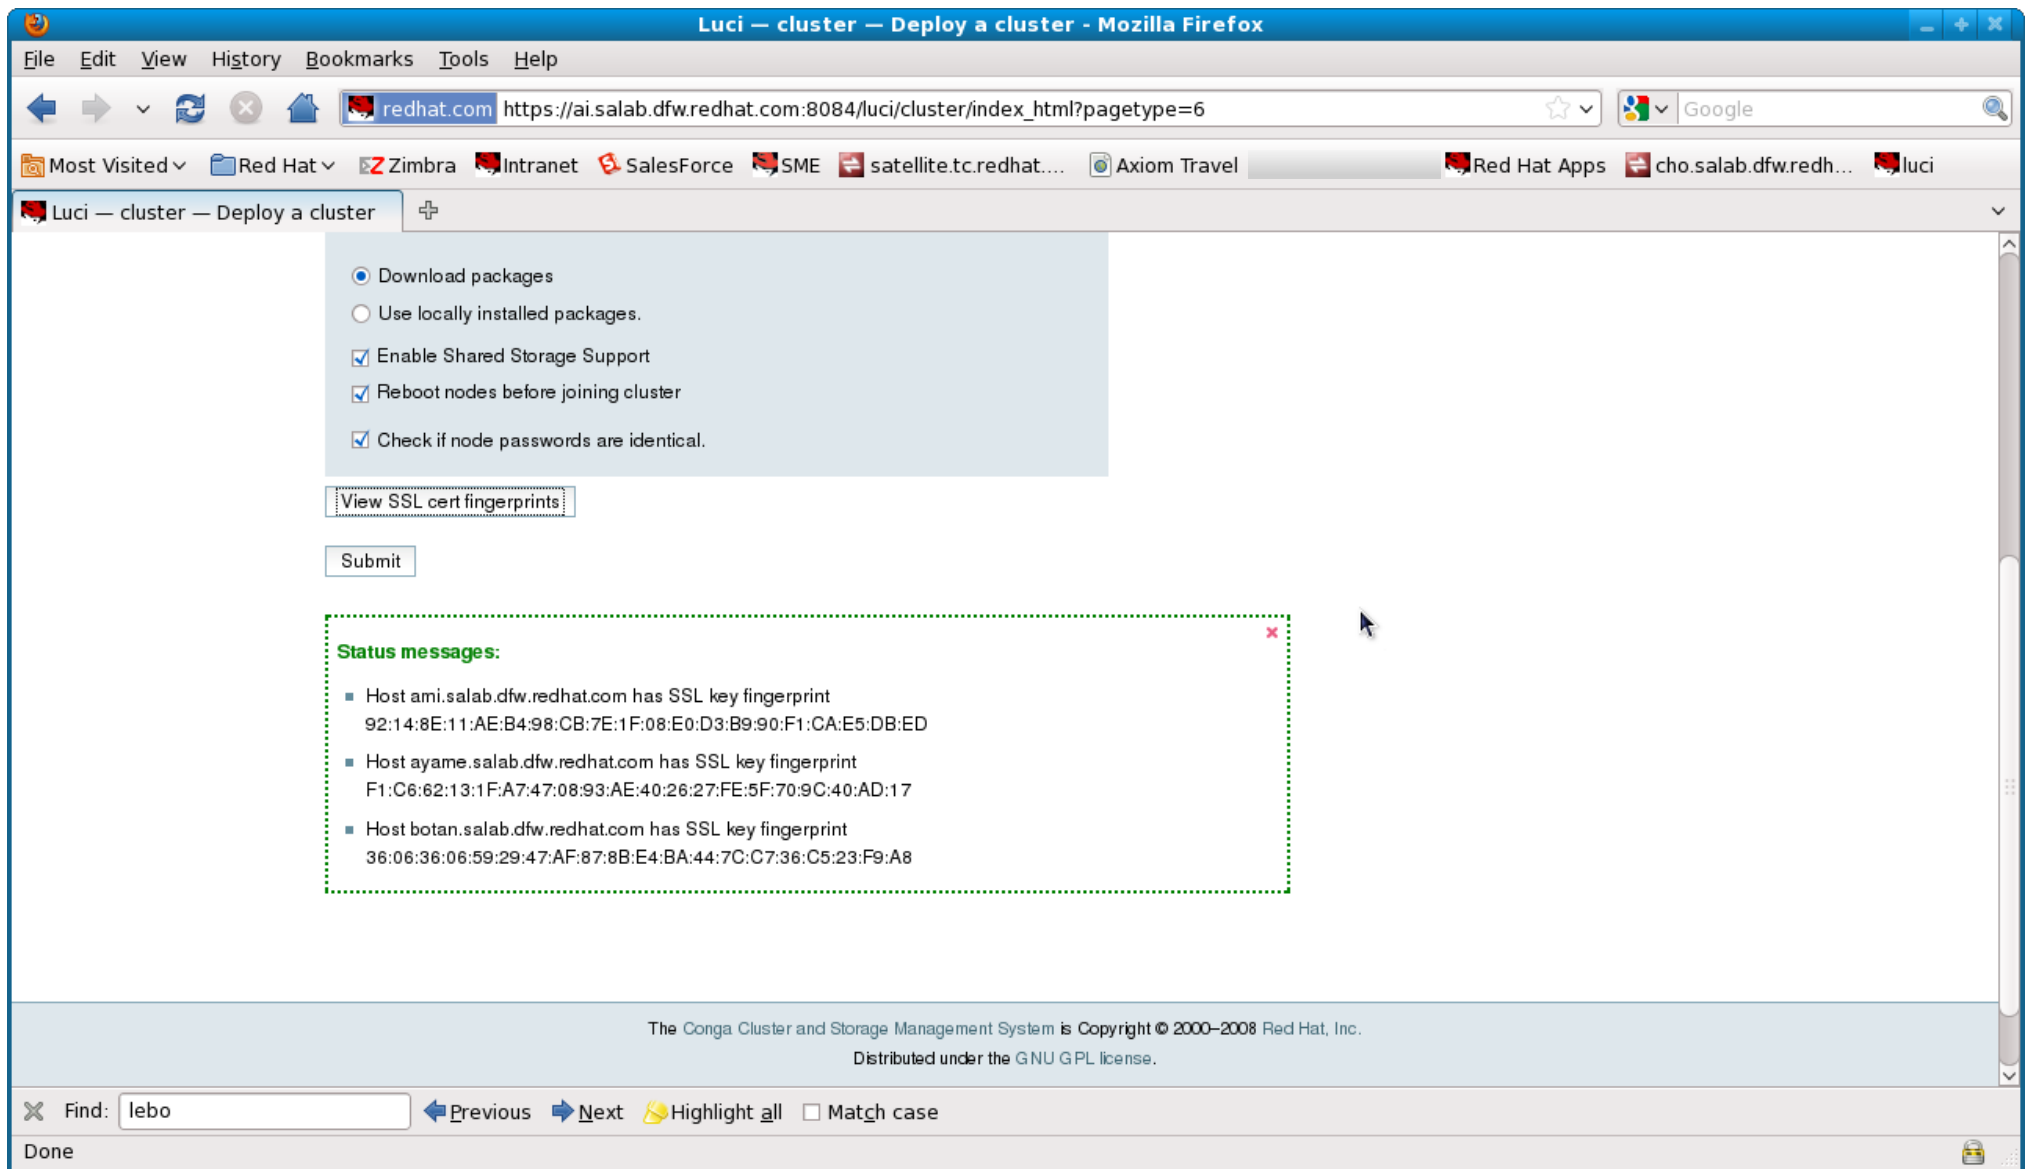
Done

Configure the Cluster Using Conga

Check to see if the luci server can talk to the ricci services on all nodes

Click “View SSL fingerprints” button

If that checks out, click on “Submit”



cluster — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=dom0-cluster&busy

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

cluster — luci

redhat CLUSTER AND STORAGE SYSTEMS


homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure
- dom0-cluster

dom0-cluster


Please be patient - this cluster's configuration is being modified.



Creating node "ami.salab.dfw.redhat.com" for cluster "dom0-cluster"

Node still being created

Install Reboot Configure Join




[Stop waiting for this job to complete](#)

Creating node "ayame.salab.dfw.redhat.com" for cluster "dom0-cluster"

Node still being created

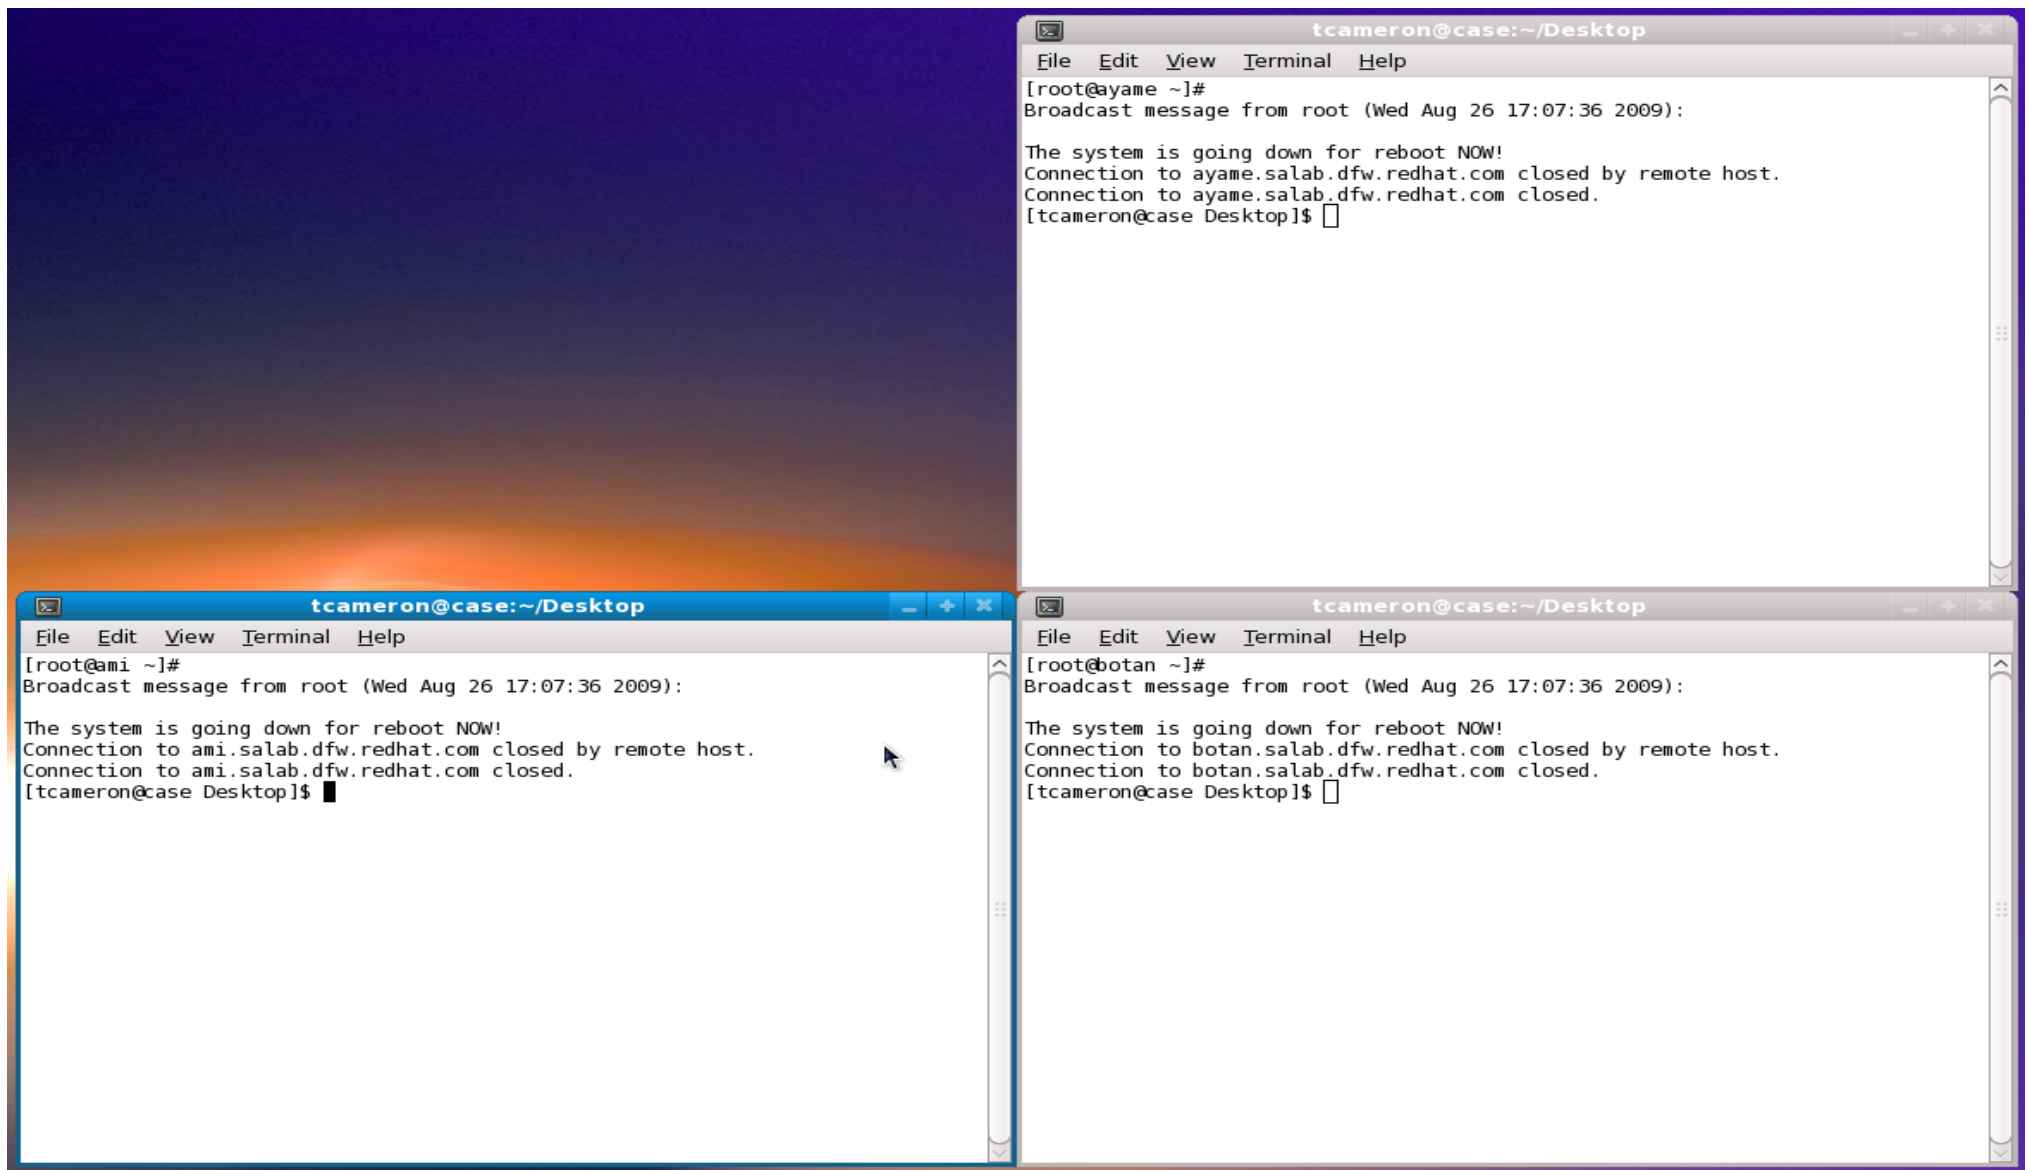
Install Reboot Configure Join

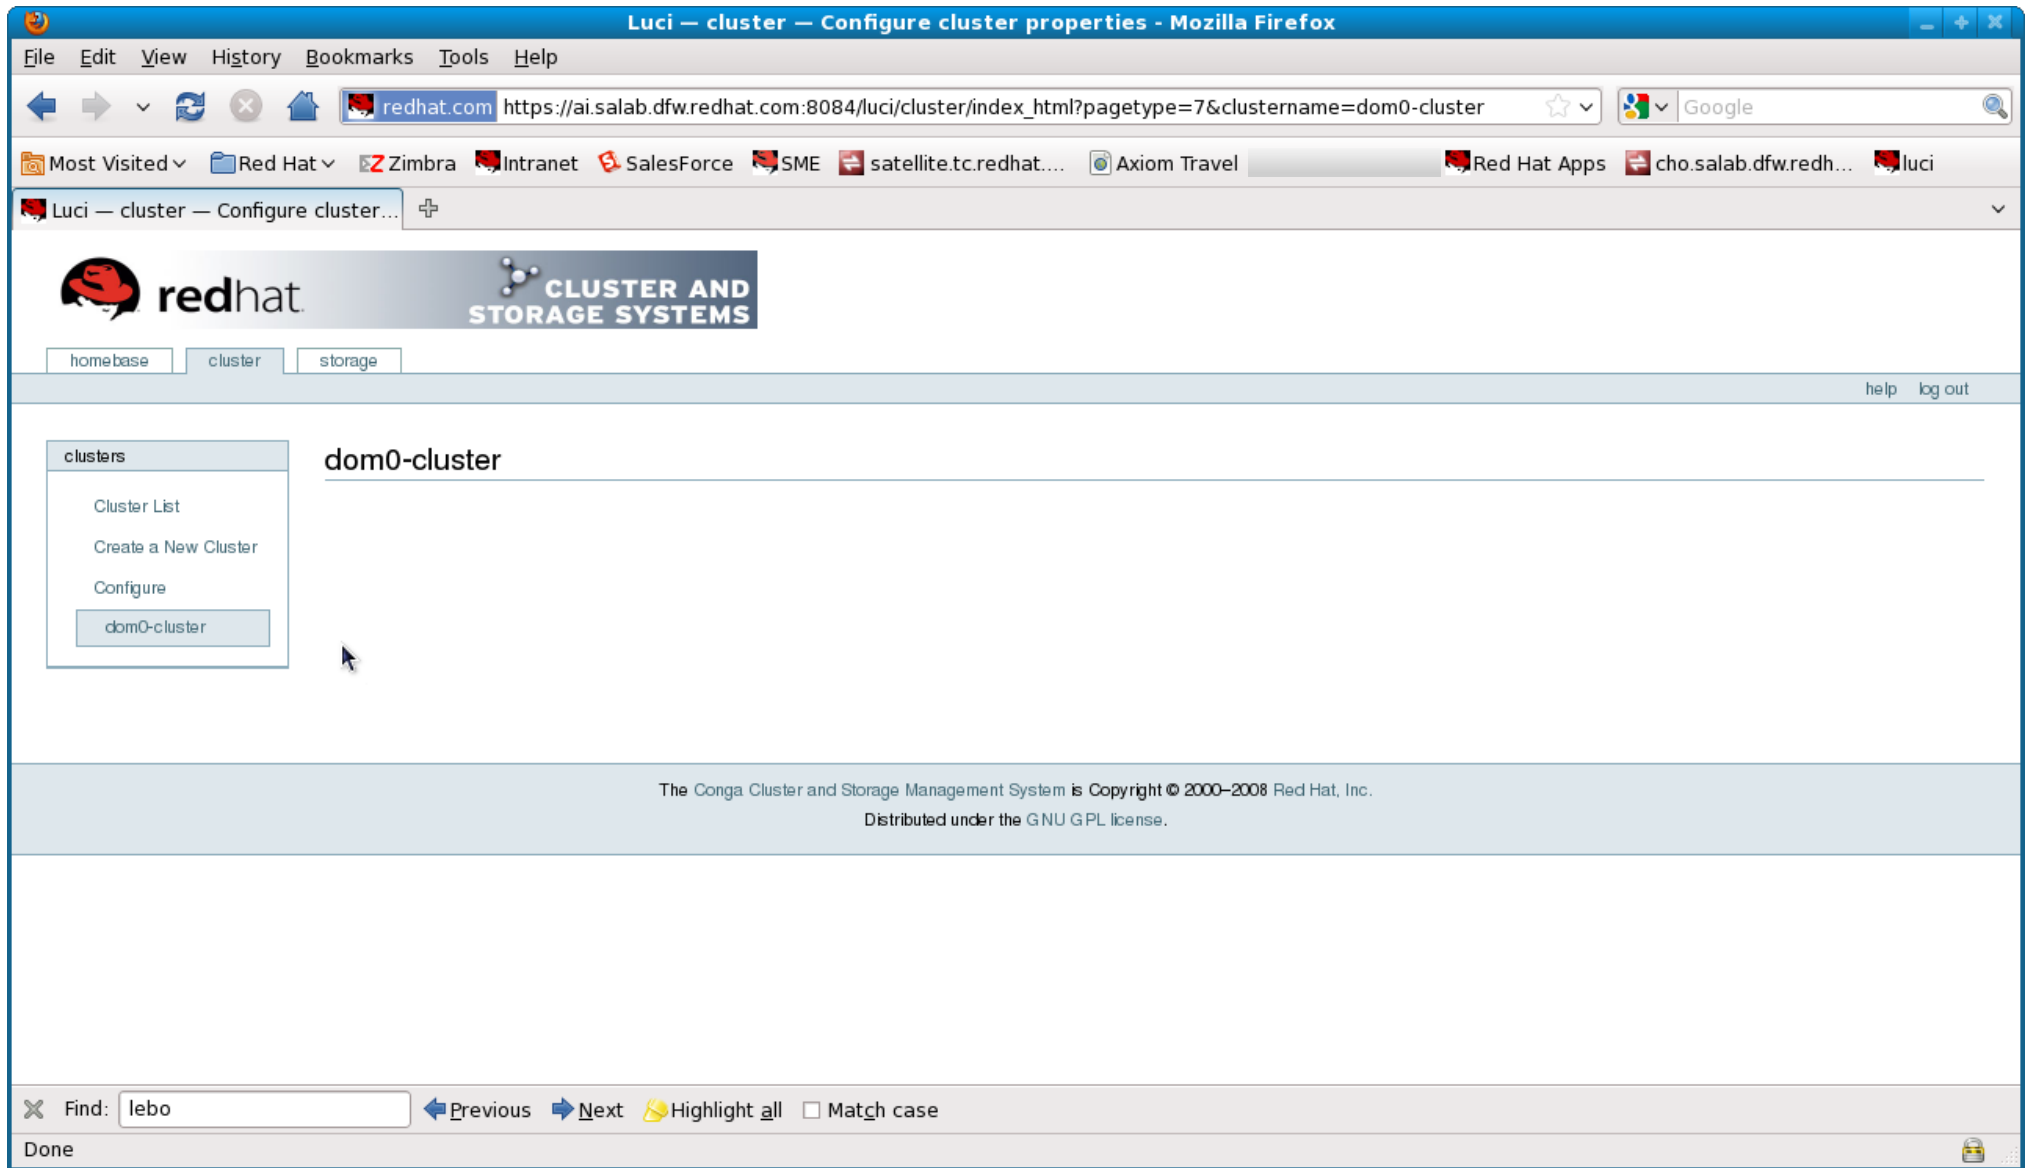


[Stop waiting for this job to complete](#)

Find: lebo Previous Next Highlight all Match case

Done





Configure the Cluster Using Conga

After the nodes all reboot, click on “Cluster List” to see if everything came up

If the cluster name and the node names are **green**, you're good.

Luci — cluster — cluster list - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=3

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — cluster list

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

Cluster List

Create a New Cluster

Configure

Choose a cluster to administer

Cluster Name: dom0-cluster Restart this cluster Go

- Status: Quorate
- Total Cluster Votes: 3
- Minimum Required Quorum: 2

Nodes

- [ami.salab.dfw.redhat.com](#)
- [ayame.salab.dfw.redhat.com](#)
- [botan.salab.dfw.redhat.com](#)

Services

- No Services Defined

The Conga Cluster and Storage Management System is Copyright © 2000–2008 Red Hat, Inc.
Distributed under the GNU GPL license.

Find: lebo Previous Next Highlight all Match case

Done

Configure the Cluster Using Conga

Set up a shared fence device

- Cluster tab

- Cluster name

- Shared Fence Devices

- Add a shareable fence device

- Fill out details

- Submit and accept

Luci — cluster — cluster list - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=3

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — cluster list

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

Cluster List

Create a New Cluster

Configure

Choose a cluster to administer

Cluster Name: dom0-cluster Restart this cluster Go

- Status: Quorate
- Total Cluster Votes: 3
- Minimum Required Quorum: 2

Nodes

- ami.salab.dfw.redhat.com
- ayame.salab.dfw.redhat.com
- botan.salab.dfw.redhat.com

Services

- No Services Defined

The Conga Cluster and Storage Management System is Copyright © 2000–2008 Red Hat, Inc.
Distributed under the GNU GPL license.

Find: lebo Previous Next Highlight all Match case

https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=dom0-cluster

Luci — cluster — Configure cluster properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — Configure cluster...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure
- dom0-cluster

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence
- Devices

dom0-cluster

General Fence Multicast Quorum Partition

General Properties

Cluster Name

Configuration Version

► Show advanced cluster properties

Find: lebo Previous Next Highlight all Match case

https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=50&clustername=dom0-cluster

Luci — cluster — shared fence devices - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=50&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — shared fence de...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence Devices
- Add a Fence Device
- Configure a Fence Device

dom0-cluster

Shared Fence Devices for Cluster: dom0-cluster

Find: lebo Previous Next Highlight all Match case

https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=51&clustername=dom0-cluster

Luci — cluster — fence devices - Add a new fence device - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=51&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — fence devices - ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence Devices
- Add a Fence Device
- Configure a Fence Device

dom0-cluster

Add a Sharable Fence Device

Fencing Type

WTI Power Switch

Fence Type WTI Power Switch

Name wti-switch

IP Address 10.15.183.249

Password

Password Script (optional)

Add this shared fence device

Find: lebo Previous Next Highlight all Match case

Done

Configure the Cluster Using Conga

Define virtual machine fencing

Cluster tab

Click on cluster

Choose Fence tab

Check the box for XVM fence daemon

Enter a hostname from dom0 cluster

Retrieve cluster nodes

Create and distribute keys

Apply and confirm

Luci — cluster — Configure cluster properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — Configure cluster...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure
- dom0-cluster

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence Devices

dom0-cluster

General Fence Multicast Quorum Partition

General Properties

Cluster Name

Configuration Version

► Show advanced cluster properties

Find: lebo Previous Next Highlight all Match case

Done

Luci — cluster — Configure cluster properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — Configure cluster... +

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure
- dom0-cluster

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence Devices

dom0-cluster

General Fence Multicast Quorum Partition

Fence Daemon Properties

Post Fail Delay 0

Post Join Delay 3

Run XVM fence daemon ☒

XVM fence daemon key distribution

Enter a node hostname from the host cluster botan

Enter a node hostname from the hosted (virtual) cluster

Retrieve cluster nodes

Apply

Find: lebo Previous Next Highlight all Match case

Done

Luci — cluster — Configure cluster properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — Configure cluster... +

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure
- dom0-cluster

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence
- Devices

dom0-cluster

General Fence Multicast Quorum Partition

Fence Daemon Properties

Post Fail Delay

Post Join Delay

Run XVM fence daemon ☒

XVM fence daemon key distribution

Cluster: dom0-cluster

- ☒ ami.salab.dfw.redhat.com
- ☒ ayame.salab.dfw.redhat.com
- ☒ botan.salab.dfw.redhat.com

Create and distribute keys

Apply

Find: lebo Previous Next Highlight all Match case

Done

cluster — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=dom0-cluster&tab=

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

cluster — luci

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters


- Cluster List
- Create a New Cluster
- Configure
- dom0-cluster

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence
- Devices

dom0-cluster

Please be patient - this cluster's configuration is being modified.



Creating fence_xvm key file for host "ami.salab.dfw.redhat.com" -- You will be redirected in 5 seconds.
[Stop waiting for this job to complete](#)

Creating fence_xvm key file for host "ayame.salab.dfw.redhat.com" -- You will be redirected in 5 seconds.
[Stop waiting for this job to complete](#)

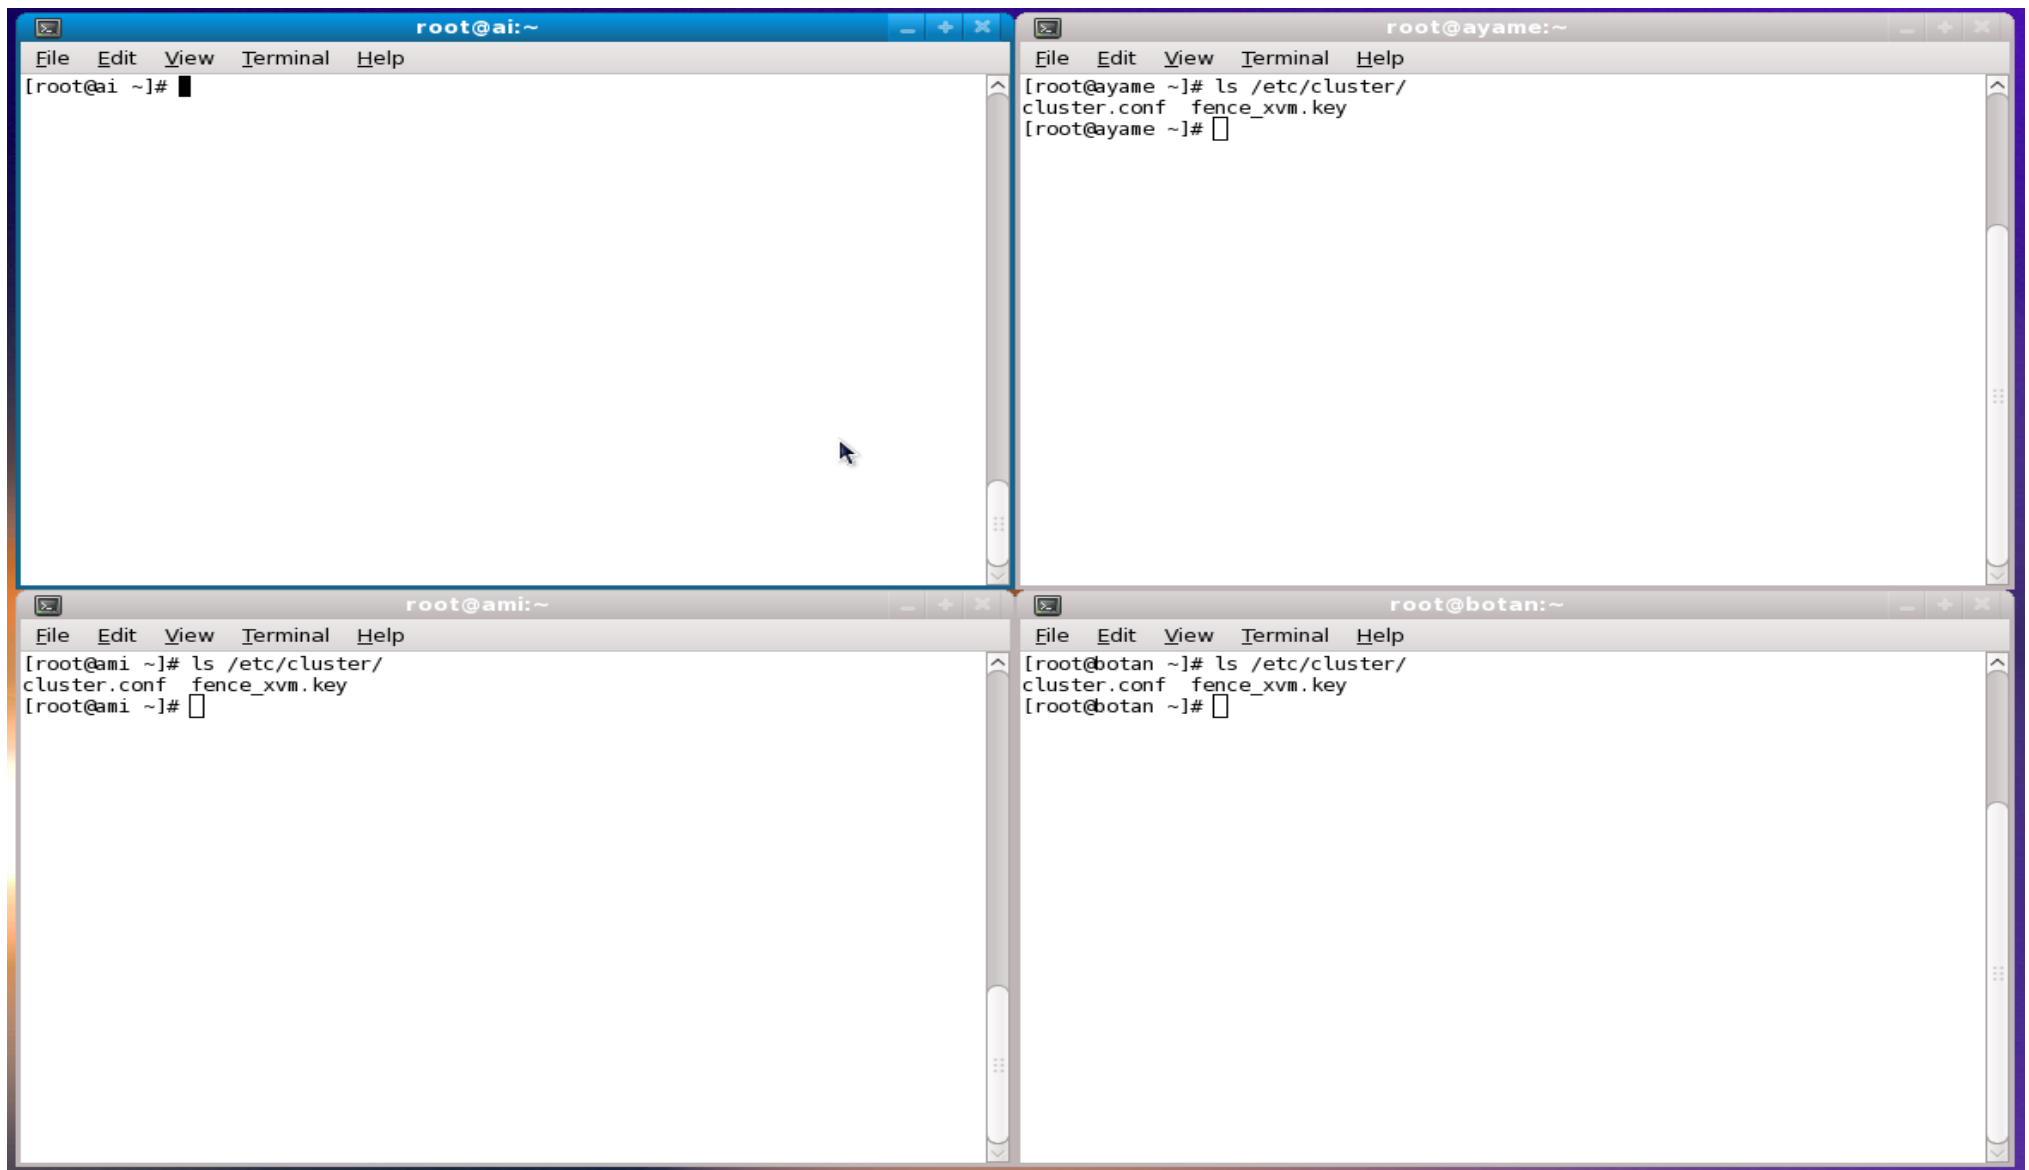
Creating fence_xvm key file for host "botan.salab.dfw.redhat.com" -- You will be redirected in 5 seconds.
[Stop waiting for this job to complete](#)

Find: lebo Previous Next Highlight all Match case

Done

Configure the Cluster Using Conga

Check that `/etc/fence_xvm.key` was created on each of the dom0 cluster node members.



Configure Per-Node Fencing

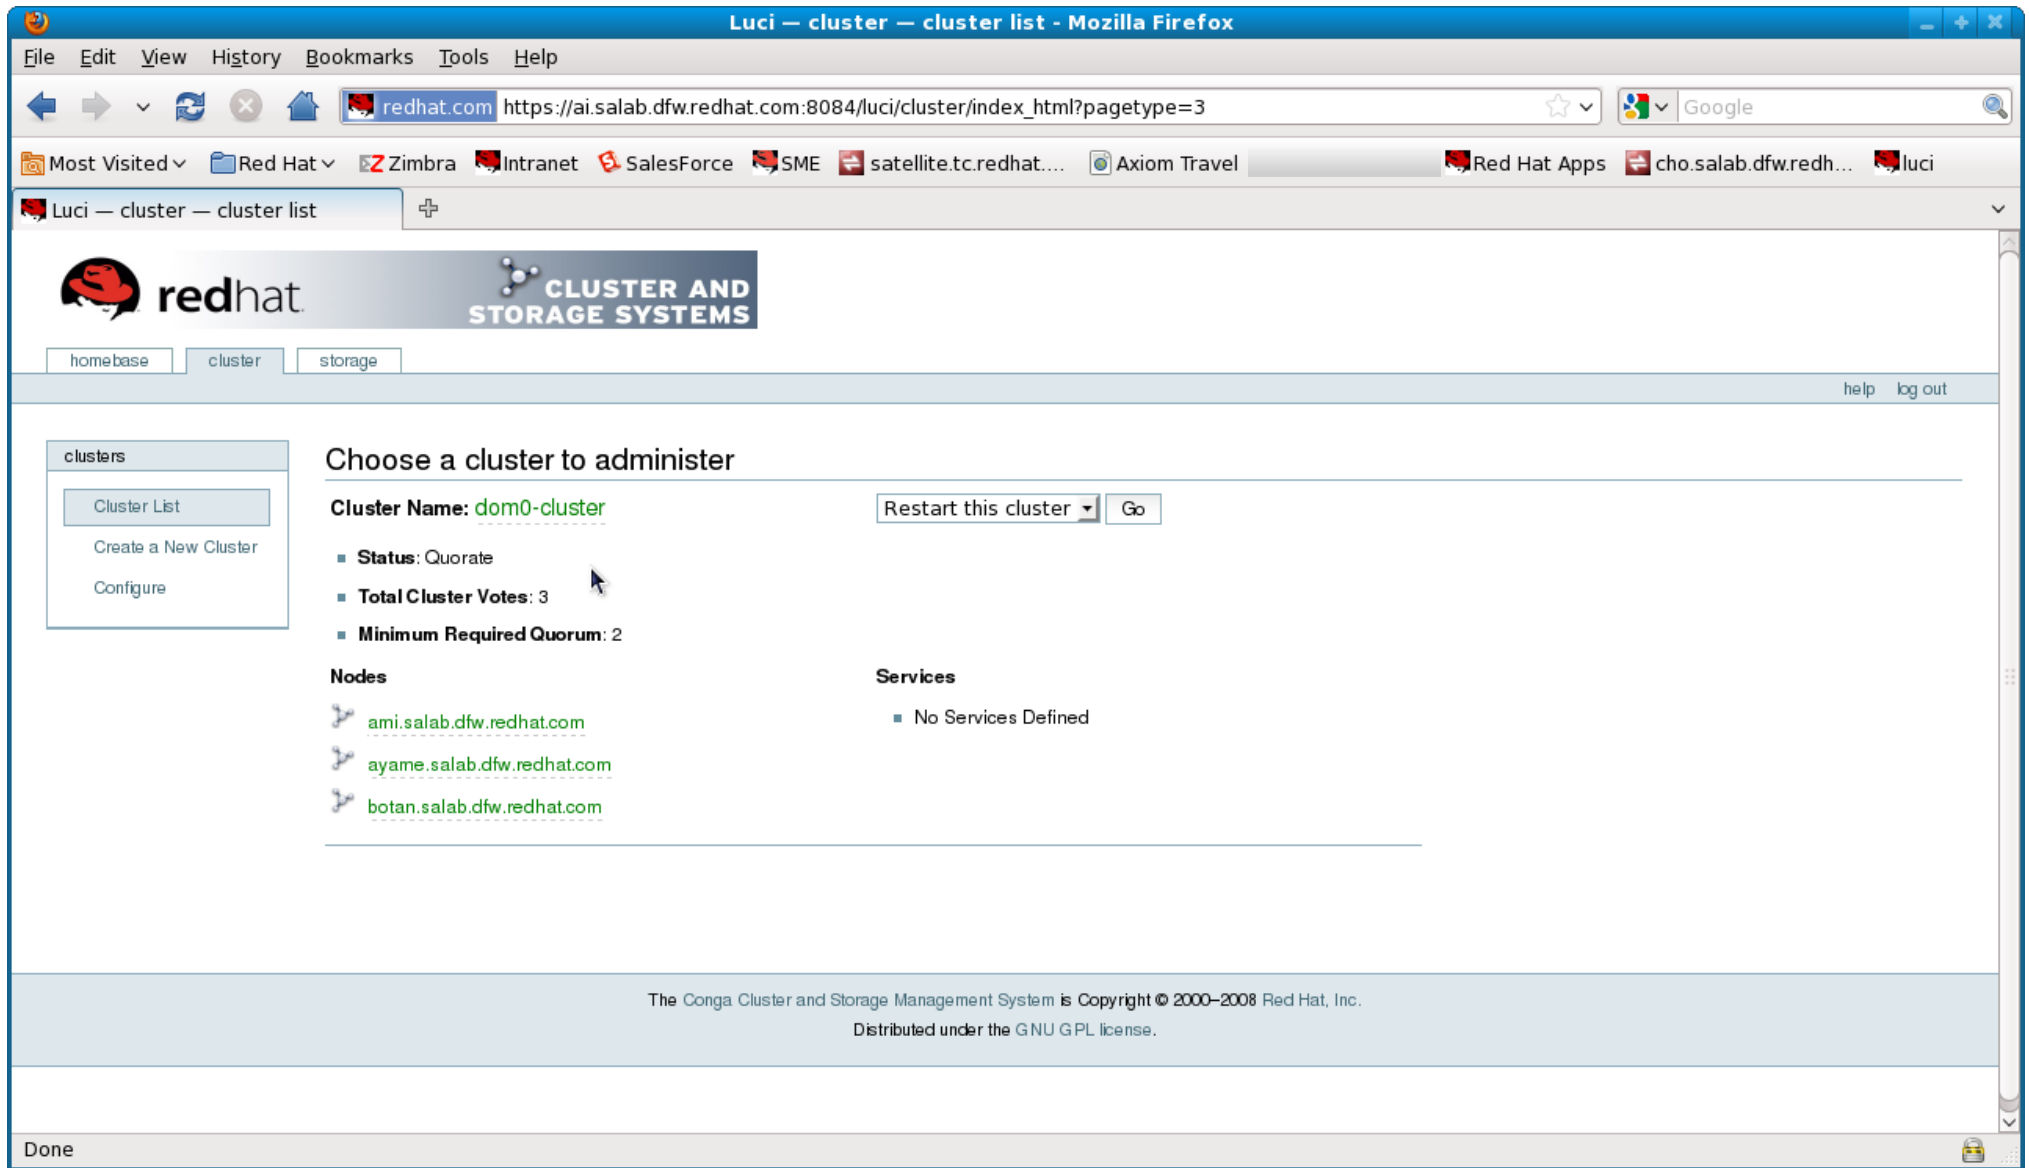
Cluster tab

Cluster Name

Node Name

Main Fencing Method

Add a fence device to this level



Luci — cluster — nodes — properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&clustername=dom0-cluster&node=ami.salab.dfw.redhat.com

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — nodes — proper...


clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
 - Add a Node
 - Configure
 - ami.salab.dfw.redhat.com
 - ayame.salab.dfw.redhat.com
 - botan.salab.dfw.redhat.com
- Services
- Resources
- Failover Domains
- Shared Fence Devices

dom0-cluster

 **Node Name:** ami.salab.dfw.redhat.com

Status: Cluster member

[Show recent log activity for this node](#)

Choose a Task... Go

Cluster daemons running on this node

| Daemon | Currently running | Enabled at start-up |
|-----------|-------------------|-------------------------------------|
| cman | yes | <input checked="" type="checkbox"/> |
| rgmanager | yes | <input checked="" type="checkbox"/> |

[Update node daemon properties](#)

Services on this Node

- No cluster services are currently running here

Failover Domain Membership

- This node has no failover domain membership

Main Fencing Method

[Add a fence device to this level](#)

Backup Fencing Method

[Add a fence device to this level](#)

Done

Luci — cluster — nodes — properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&nodename=ami.salab.dfw.redhat

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — nodes — proper...

Configure

- ami.salab.dfw.redhat.com
- ayame.salab.dfw.redhat.com
- botan.salab.dfw.redhat.com

Services

Resources

Failover Domains

Shared Fence Devices

cman yes ☒

rgmanager yes ☒

Update node daemon properties

Services on this Node

- No cluster services are currently running here

Failover Domain Membership

- This node has no failover domain membership

Main Fencing Method

----- Use an existing Fence Device -----

Add a fence device to this level

Update main fence properties

Backup Fencing Method

Add a fence device to this level

Update backup fence properties

The Conga Cluster and Storage Management System is Copyright © 2000–2008 Red Hat, Inc.
Distributed under the GNU GPL license.

Done

Luci — cluster — nodes — properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&nodename=ami.salab.dfw.redhat

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — nodes — proper... +

- Services
- Resources
- Failover Domains
- Shared Fence Devices

Services on this Node

- No cluster services are currently running here

Failover Domain Membership

- This node has no failover domain membership

Main Fencing Method

Fence Type WTI Power Switch

Name

IP Address

Password

Password Script (optional)

Port

Use SSH ☐

[Add a fence device to this level](#)

Backup Fencing Method

[Add a fence device to this level](#)

Done

Configure Per-Node Fencing

Lather, rinse, repeat for each node.

Configure Storage Using Conga

Set up the volume group which will be used to house the virtual machine

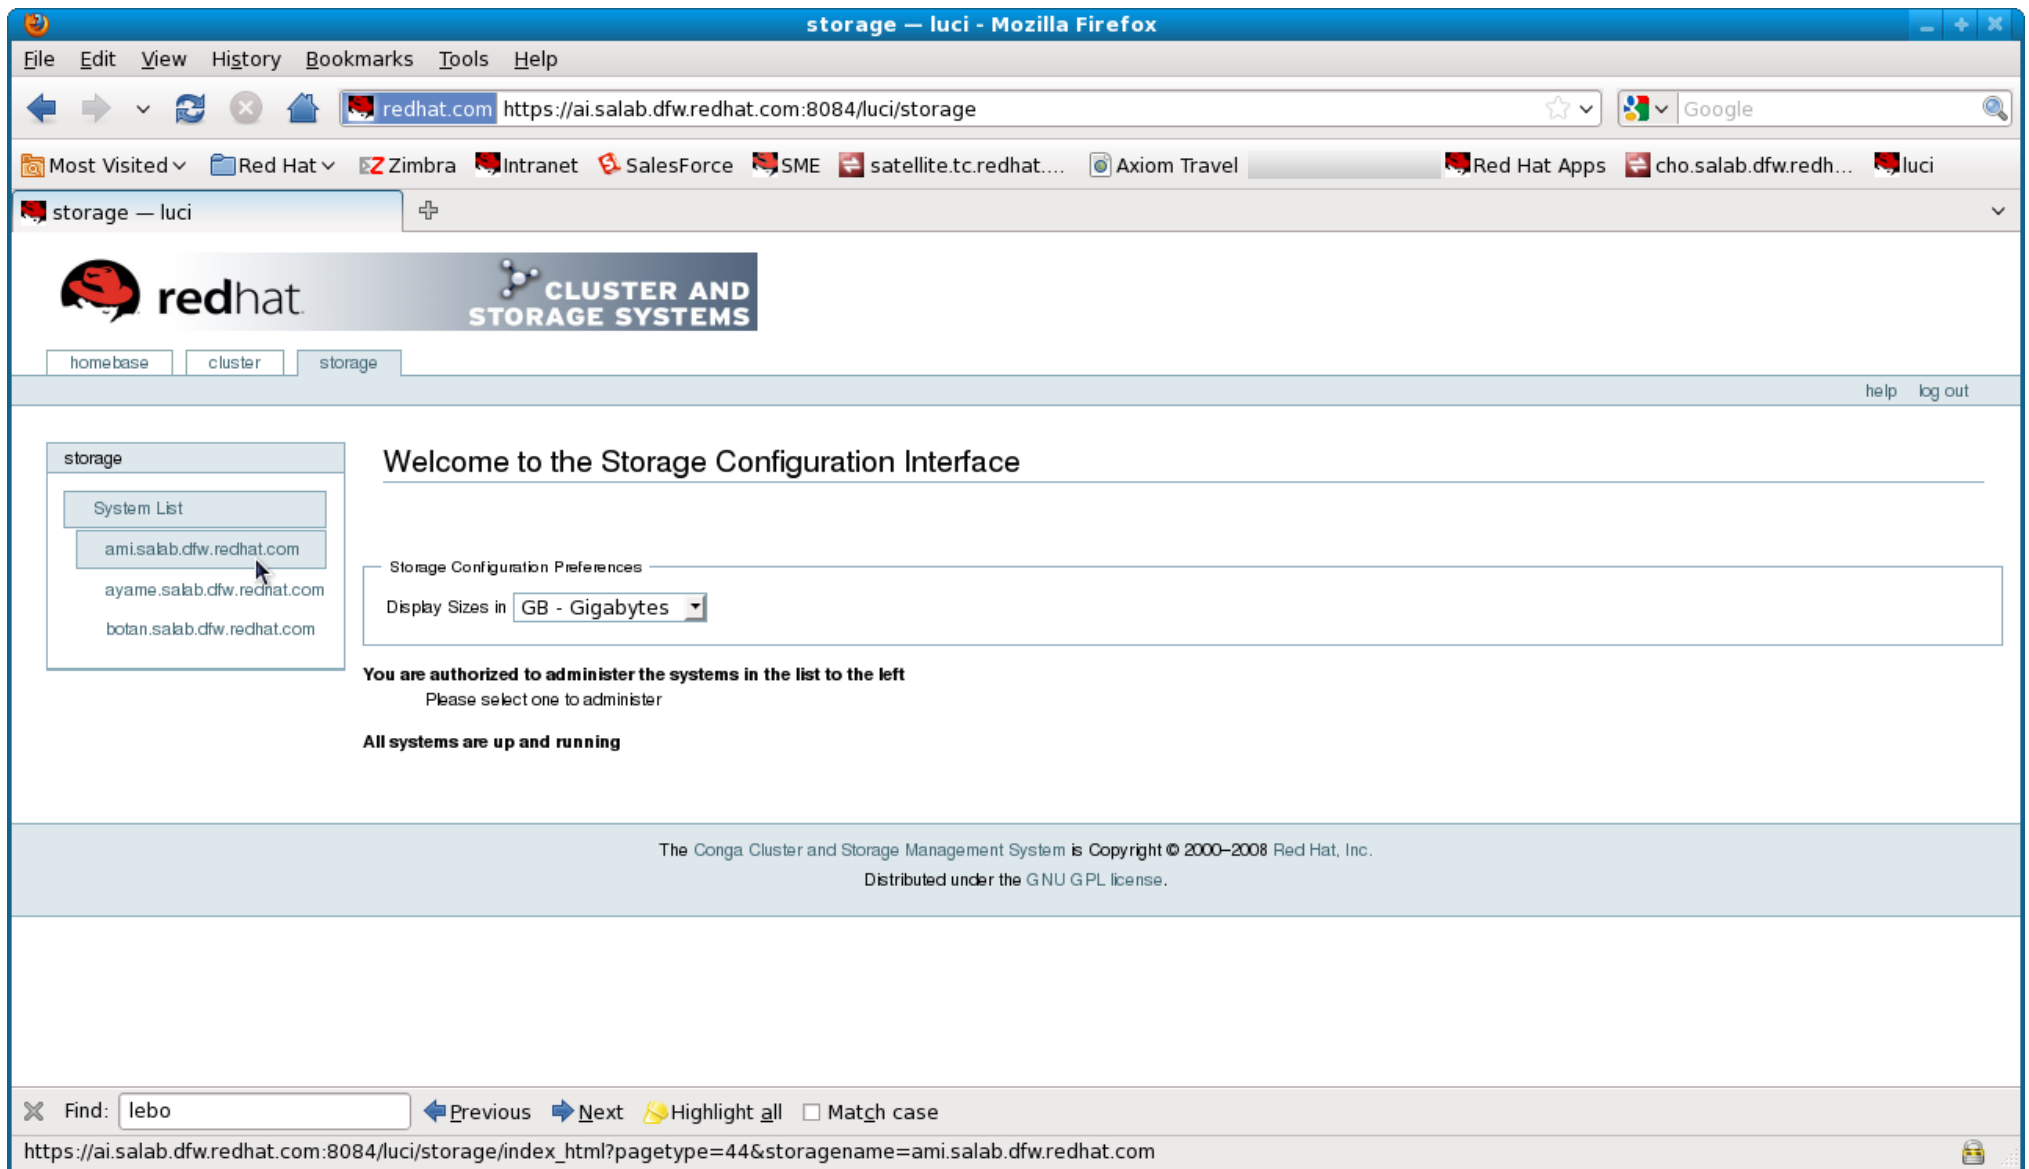
Storage tab

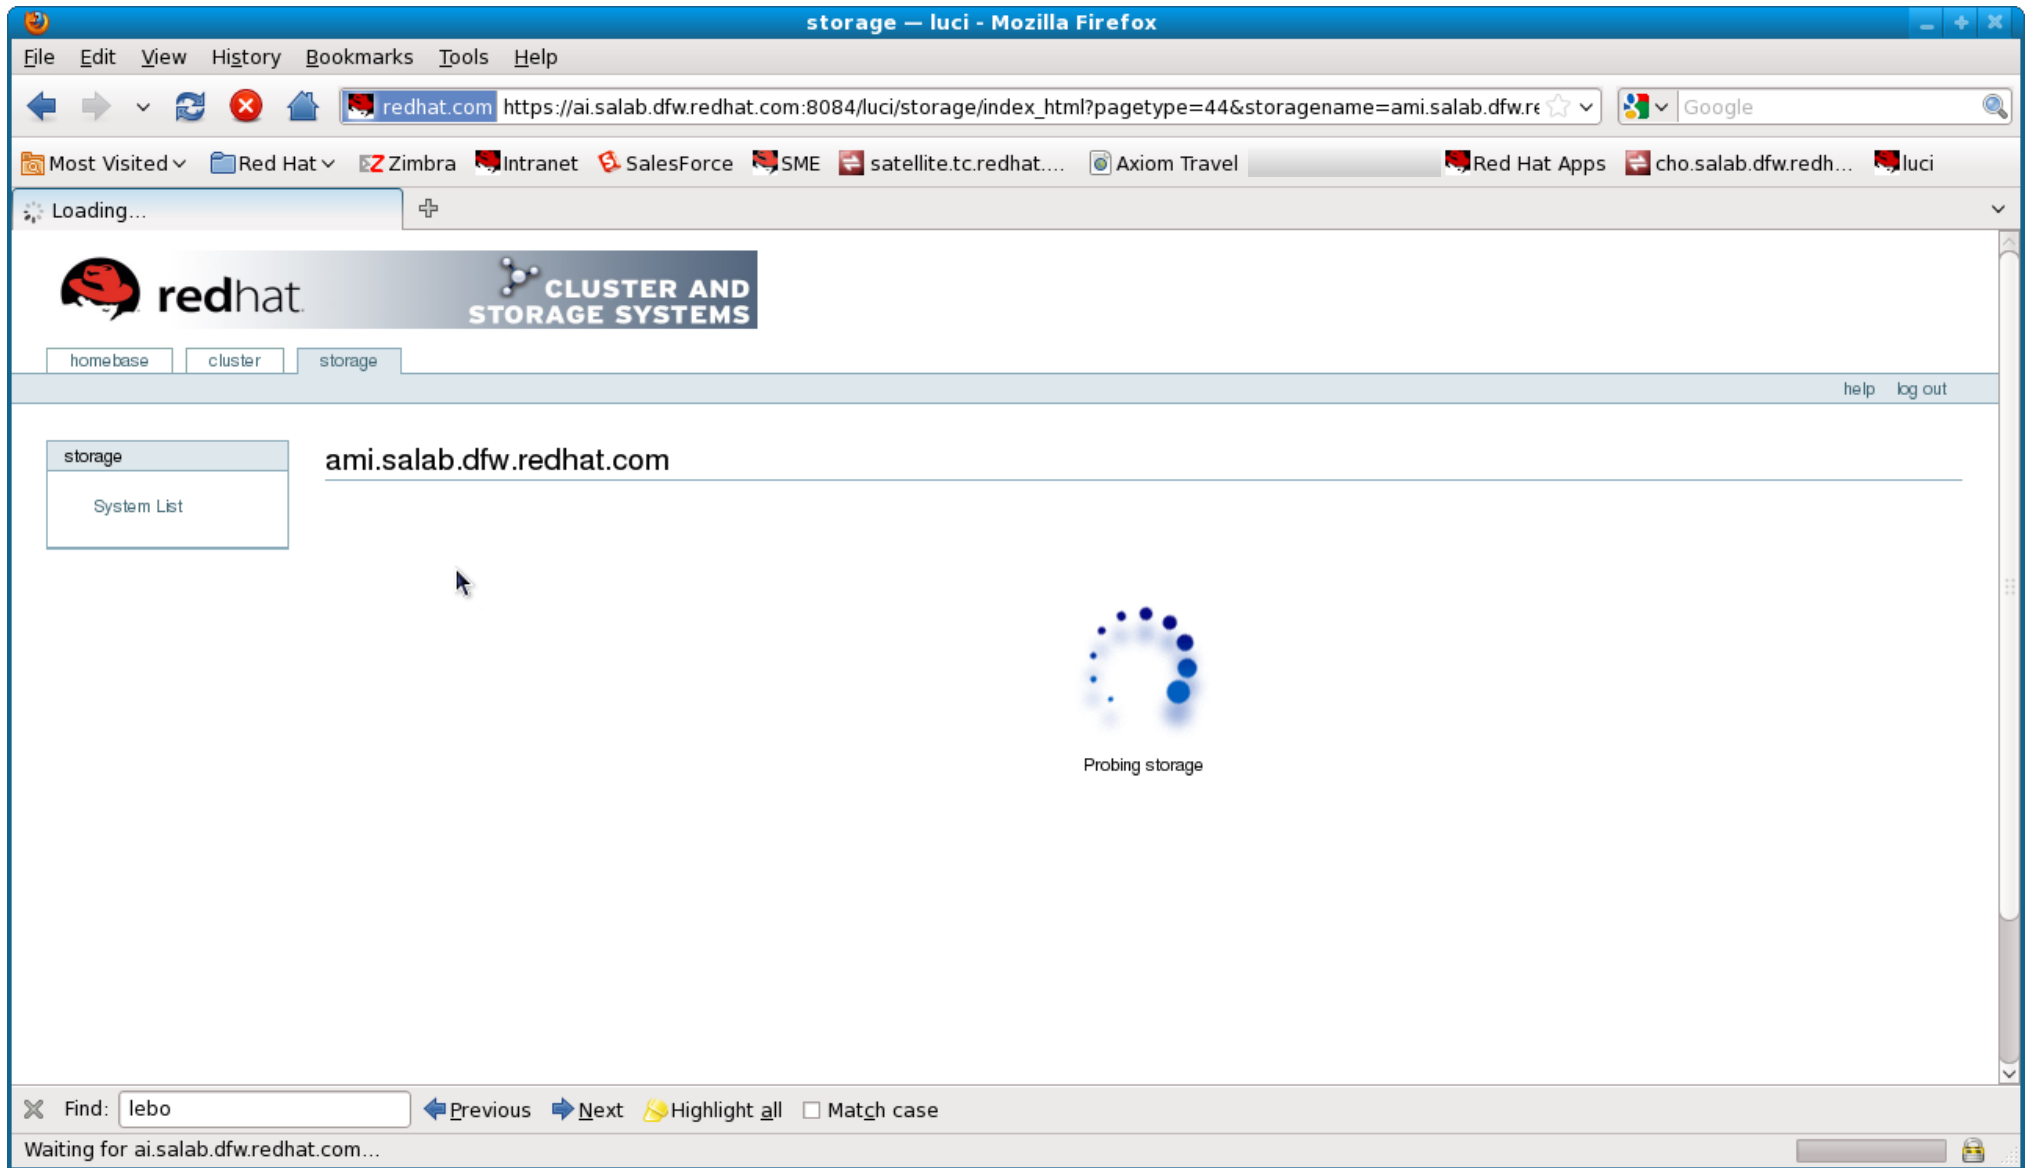
Choose a server

Choose Volume Groups

Choose New Volume Group

Name new VG and assign to a block device





storage — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/storage/index_html?pagetype=44&storagename=ami.salab.dfw.re

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

storage — luci

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

storage

System List

ami.salab.dfw.redhat.com

Hard Drives

- sda 931.51 GB, SCSI ID = SATA_ST31000333AS_9TE1JWJ4
- sdb 100.0 GB, SCSI ID = S_beaf11

Partition Tables

- sda
- sdb

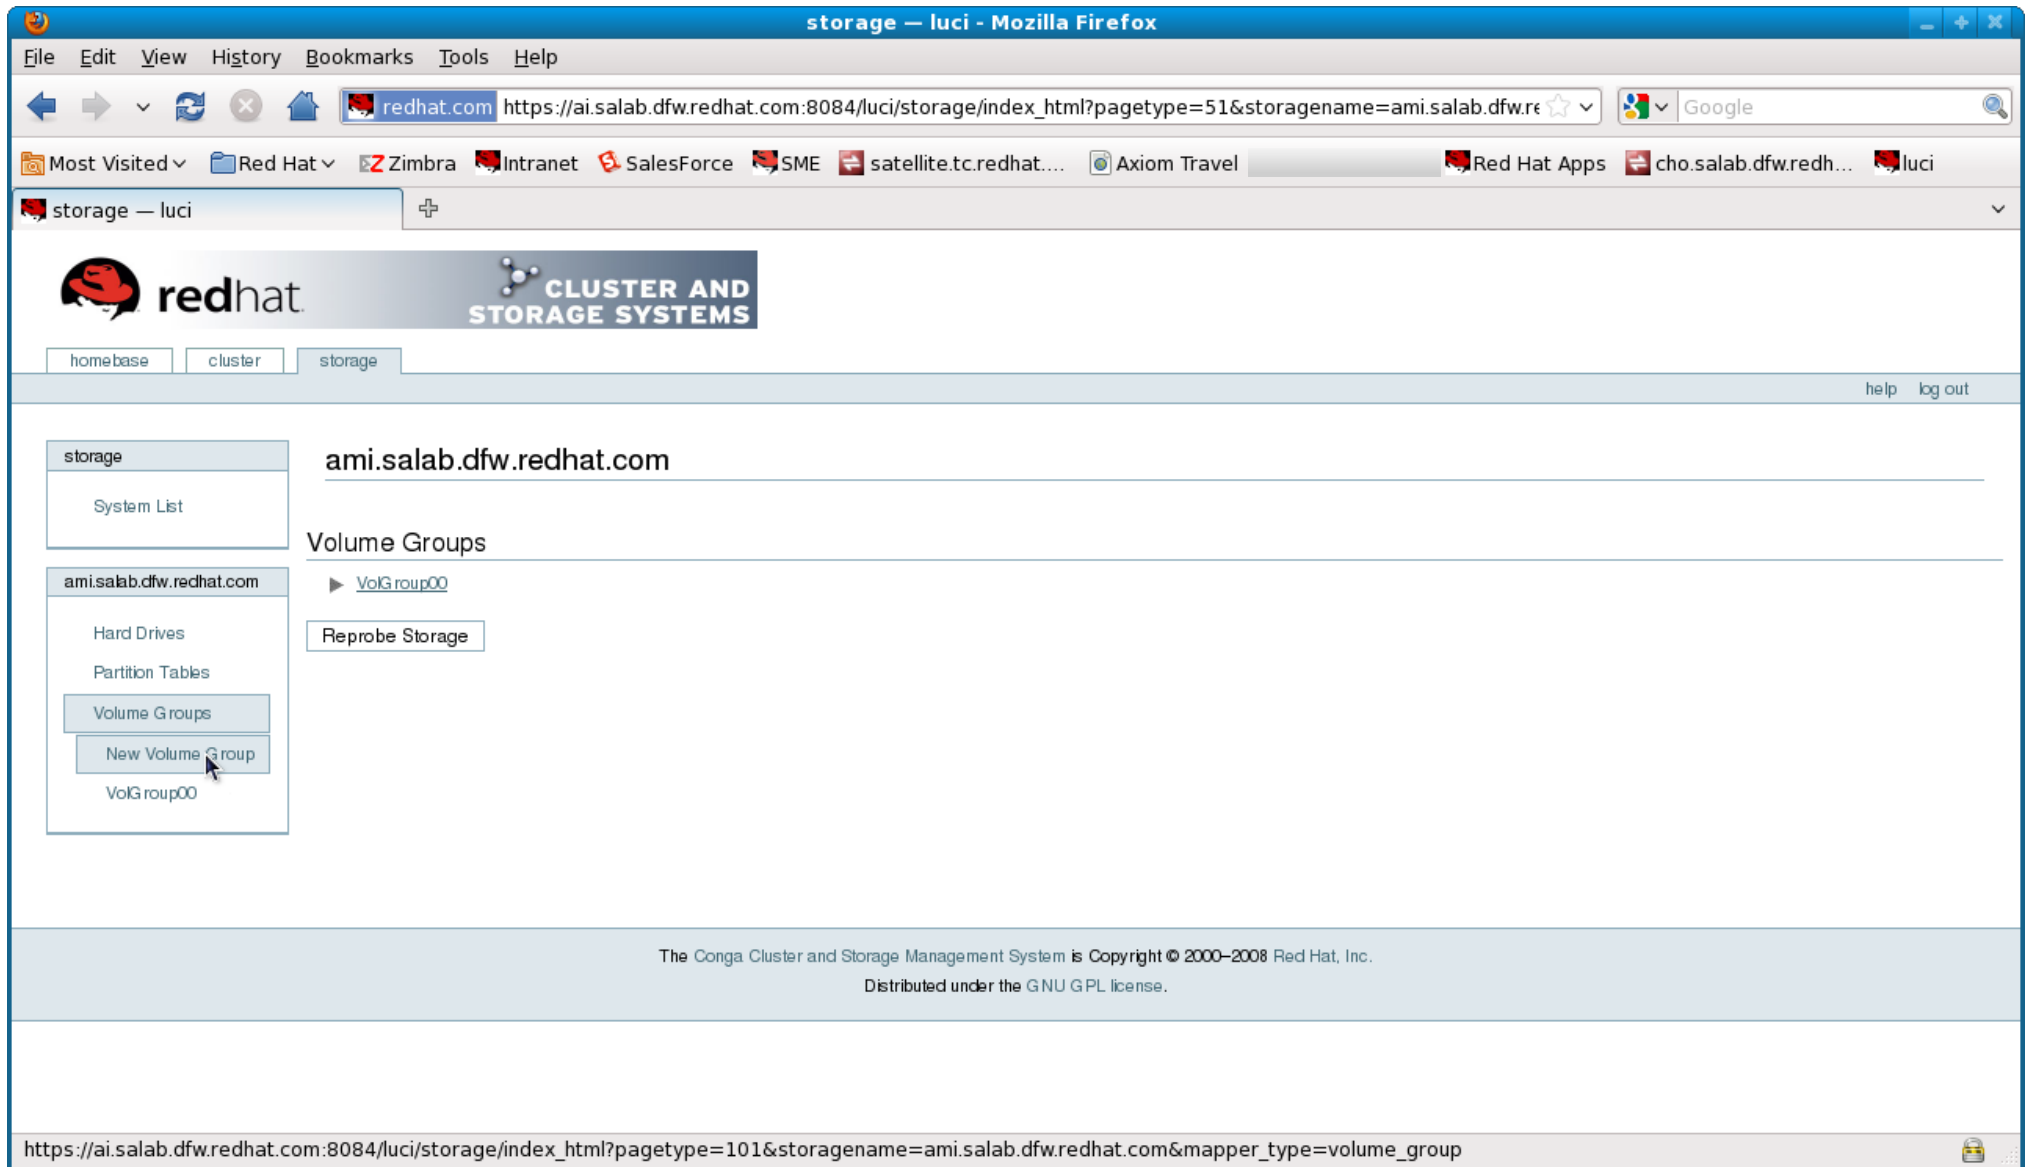
Volume Groups

- VolGroup00

Reprobe Storage

The Conga Cluster and Storage Management System is Copyright © 2000–2008 Red Hat, Inc.
Distributed under the GNU GPL license.

https://ai.salab.dfw.redhat.com:8084/luci/storage/index_html?pagetype=51&storagename=ami.salab.dfw.redhat.com&mapper_type=volume_group



storage — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/storage/index_html?pagetype=101&storagename=ami.salab.dfw...

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

storage — luci

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

storage

System List

ami.salab.dfw.redhat.com

Hard Drives

Partition Tables

Volume Groups

New Volume Group

VolGroup00

ami.salab.dfw.redhat.com

Creating New Volume Group

Volume Group Name XenVG

Extent Size 4.0 MB

Clustered true

Select 1 Physical Volume

☒ /dev/sdb1 (99.99 GB - Partition)

Reset Create

Reprobe Storage

The Conga Cluster and Storage Management System is Copyright © 2000–2008 Red Hat, Inc.
Distributed under the GNU GPL license.

Done

storage — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/storage/index_html?storagename=ami.salab.dfw.redhat.com&map

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

storage — luci

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

storage

System List

ami.salab.dfw.redhat.com

Hard Drives

Partition Tables

Volume Groups

VolGroup00

XenVG

ami.salab.dfw.redhat.com

Volume Group XenVG

☒ Graphical View (Uncheck if volumes are too small to select)

Logical Volumes:

Physical Volumes:

Click cylinders to view properties, unselect all to view Volume Group's properties

Volume Group 'XenVG'

| | |
|-------------------|----------|
| Volume Group Name | XenVG |
| Extent Size | 4.0 MB |
| Total Extents | 25599 |
| Free Extents | 25599 |
| Unused | 99.99 GB |

Done

Configure Storage Using Conga

Set up the logical volume which will be used to house the virtual machines

Scroll down

New Logical Volume

Define LV Name

Define LV Size

Define Content (GFS2 recommended)

Define Unique GFS Name (gfs0)

Define mount point (/var/lib/xen/images)

Configure Storage Using Conga

Set up the logical volume which will be used to house the virtual machines

- Set Mount to true

- Set List in /etc/fstab to true

- Set number of journals

- Verify Clustered is set to true

storage — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/storage/index_html?storagename=ami.salab.dfw.redhat.com&map Google

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

storage — luci

XenVG

Click cylinders to view properties, unselect all to view Volume Group's properties

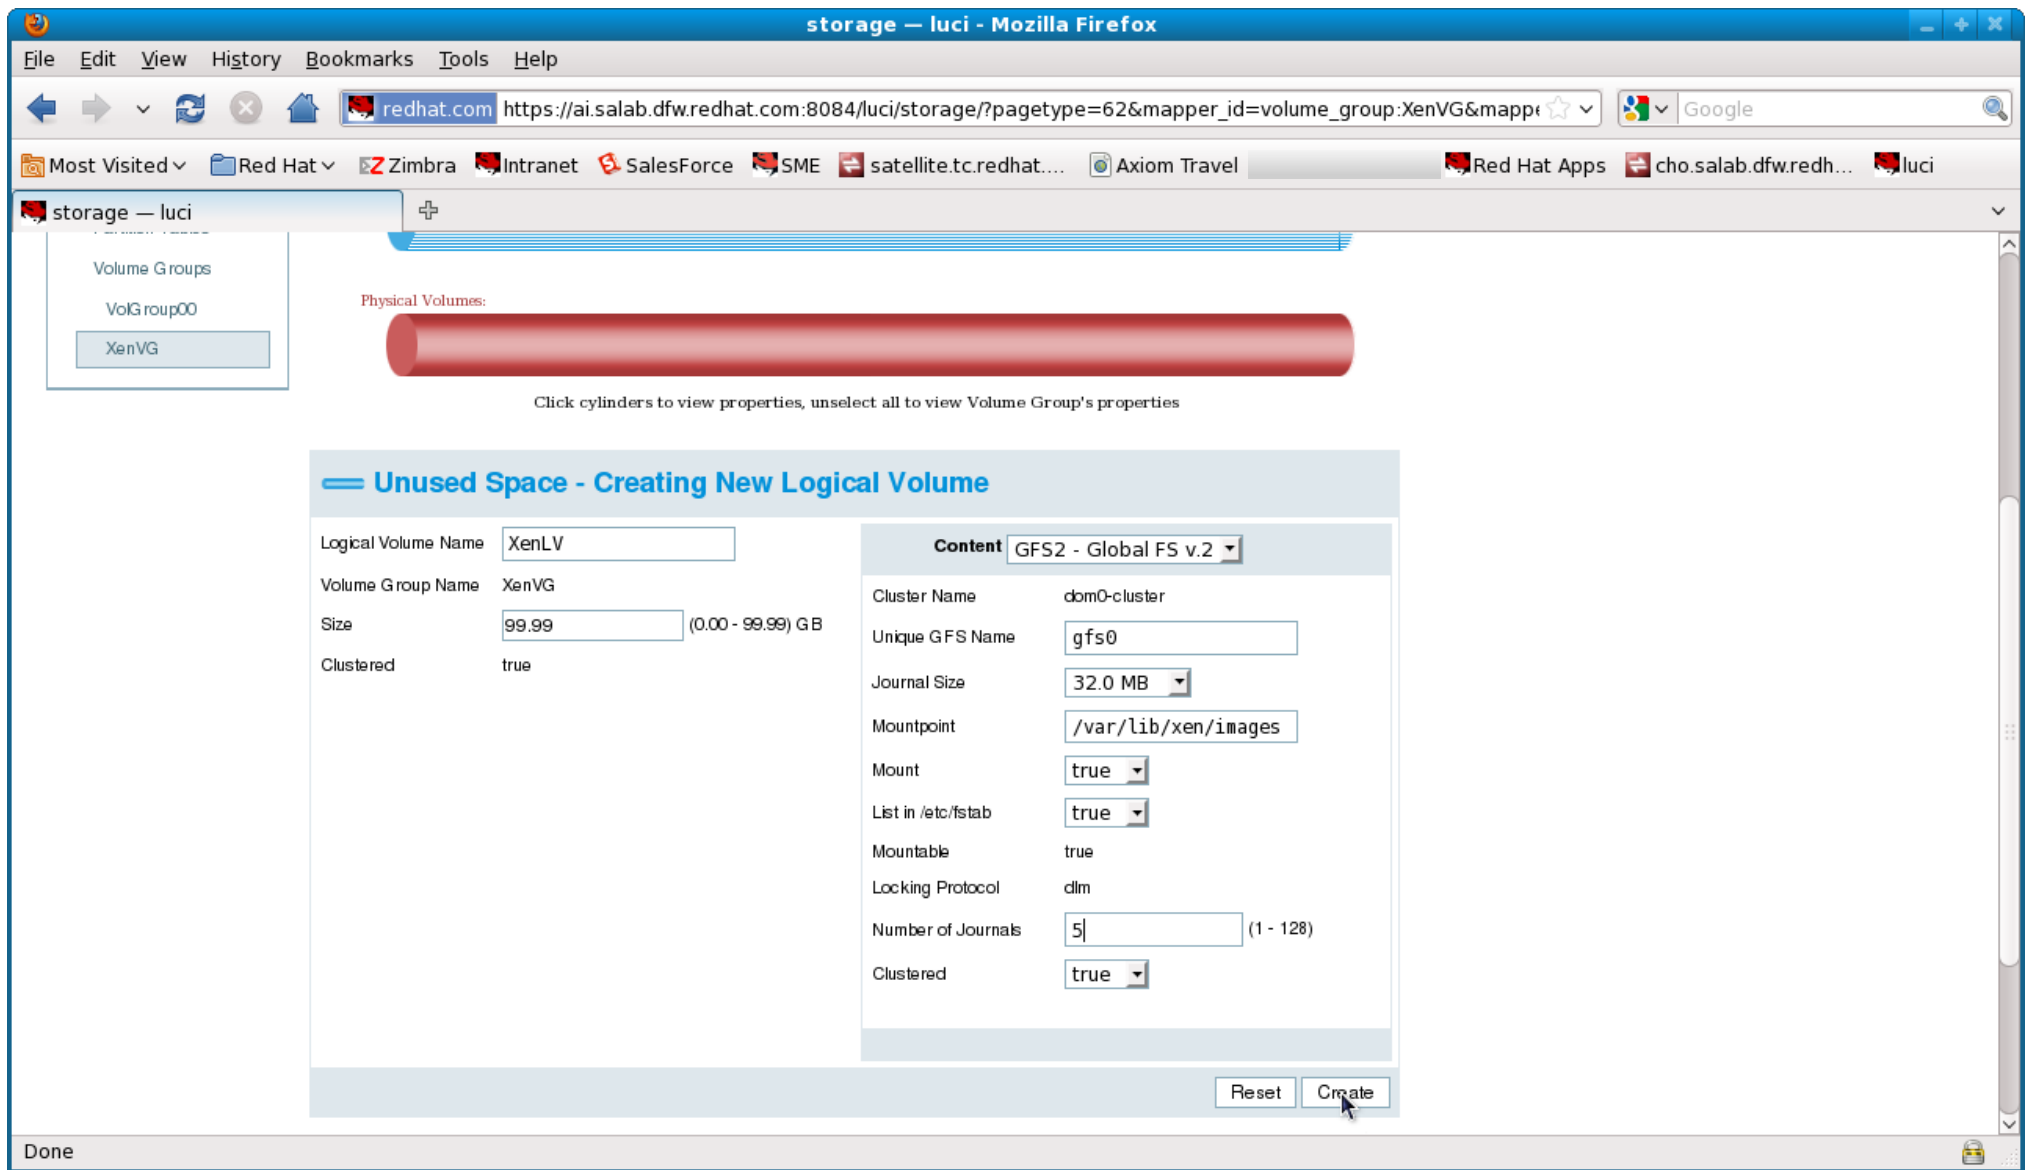
Volume Group 'XenVG'

| | |
|--------------------------|--|
| Volume Group Name | XenVG |
| Extent Size | 4.0 MB |
| Total Extents | 25599 |
| Free Extents | 25599 |
| Unused | 99.99 GB |
| Size | 99.99 GB |
| Used Extents | 0 |
| Maximum Physical Volumes | 256 |
| Maximum Logical Volumes | 256 |
| Attributes | wz--nc |
| Clustered | true |
| UUID | qcsLks-751V-vY8S-UAXL-EzP6-HHjz-H92kTy |

Remove Add Physical Volumes New Logical Volume Reset Apply

Reprobe Storage

Done



```
root@ami:~  
File Edit View Terminal Help  
15289 ? S< 0:00 [iscsi_q_5]  
15290 ? S< 0:00 [scsi_wq_5]  
19327 ? Rs 0:00 sshd: root@pts/1  
19329 pts/1 Ss 0:00 -bash  
22497 ? Ssl 0:00 /sbin/ccsd  
22507 ? SLl 0:00 aisexec  
22544 ? Ss 0:00 /sbin/groupd  
22551 ? Ss 0:00 /sbin/fenced  
22557 ? Ss 0:00 /sbin/dlm_control  
22563 ? Ss 0:00 /sbin/gfs_control  
22583 ? Ssl 0:00 clvmd -T20  
22584 ? S< 0:00 [dlm_astd]  
22585 ? S< 0:00 [dlm_scmd]  
22586 ? S< 0:00 [dlm_rcv]  
22587 ? S< 0:00 [dlm_send]  
22588 ? S< 0:00 [dlm_recoverd]  
22641 ? S<Ls 0:00 clurgmgrd  
22643 ? S<l 0:00 clurgmgrd  
22645 ? S< 0:00 [dlm_recoverd]  
23175 ? S<s 0:00 /usr/libexec/ricci/ricci-worker -f /var/lib/ricci/que  
23176 ? S 0:00 ricci-modstorage  
23266 ? R 0:00 /sbin/mkfs.gfs2 -J 32 -j 5 -0 -p lock_dlm -t dom0-clu  
23269 pts/1 R+ 0:00 ps ax  
[root@ami ~]#
```

Configure Storage Using Conga

Add the mount point to the other hosts in the cluster

Storage

System List

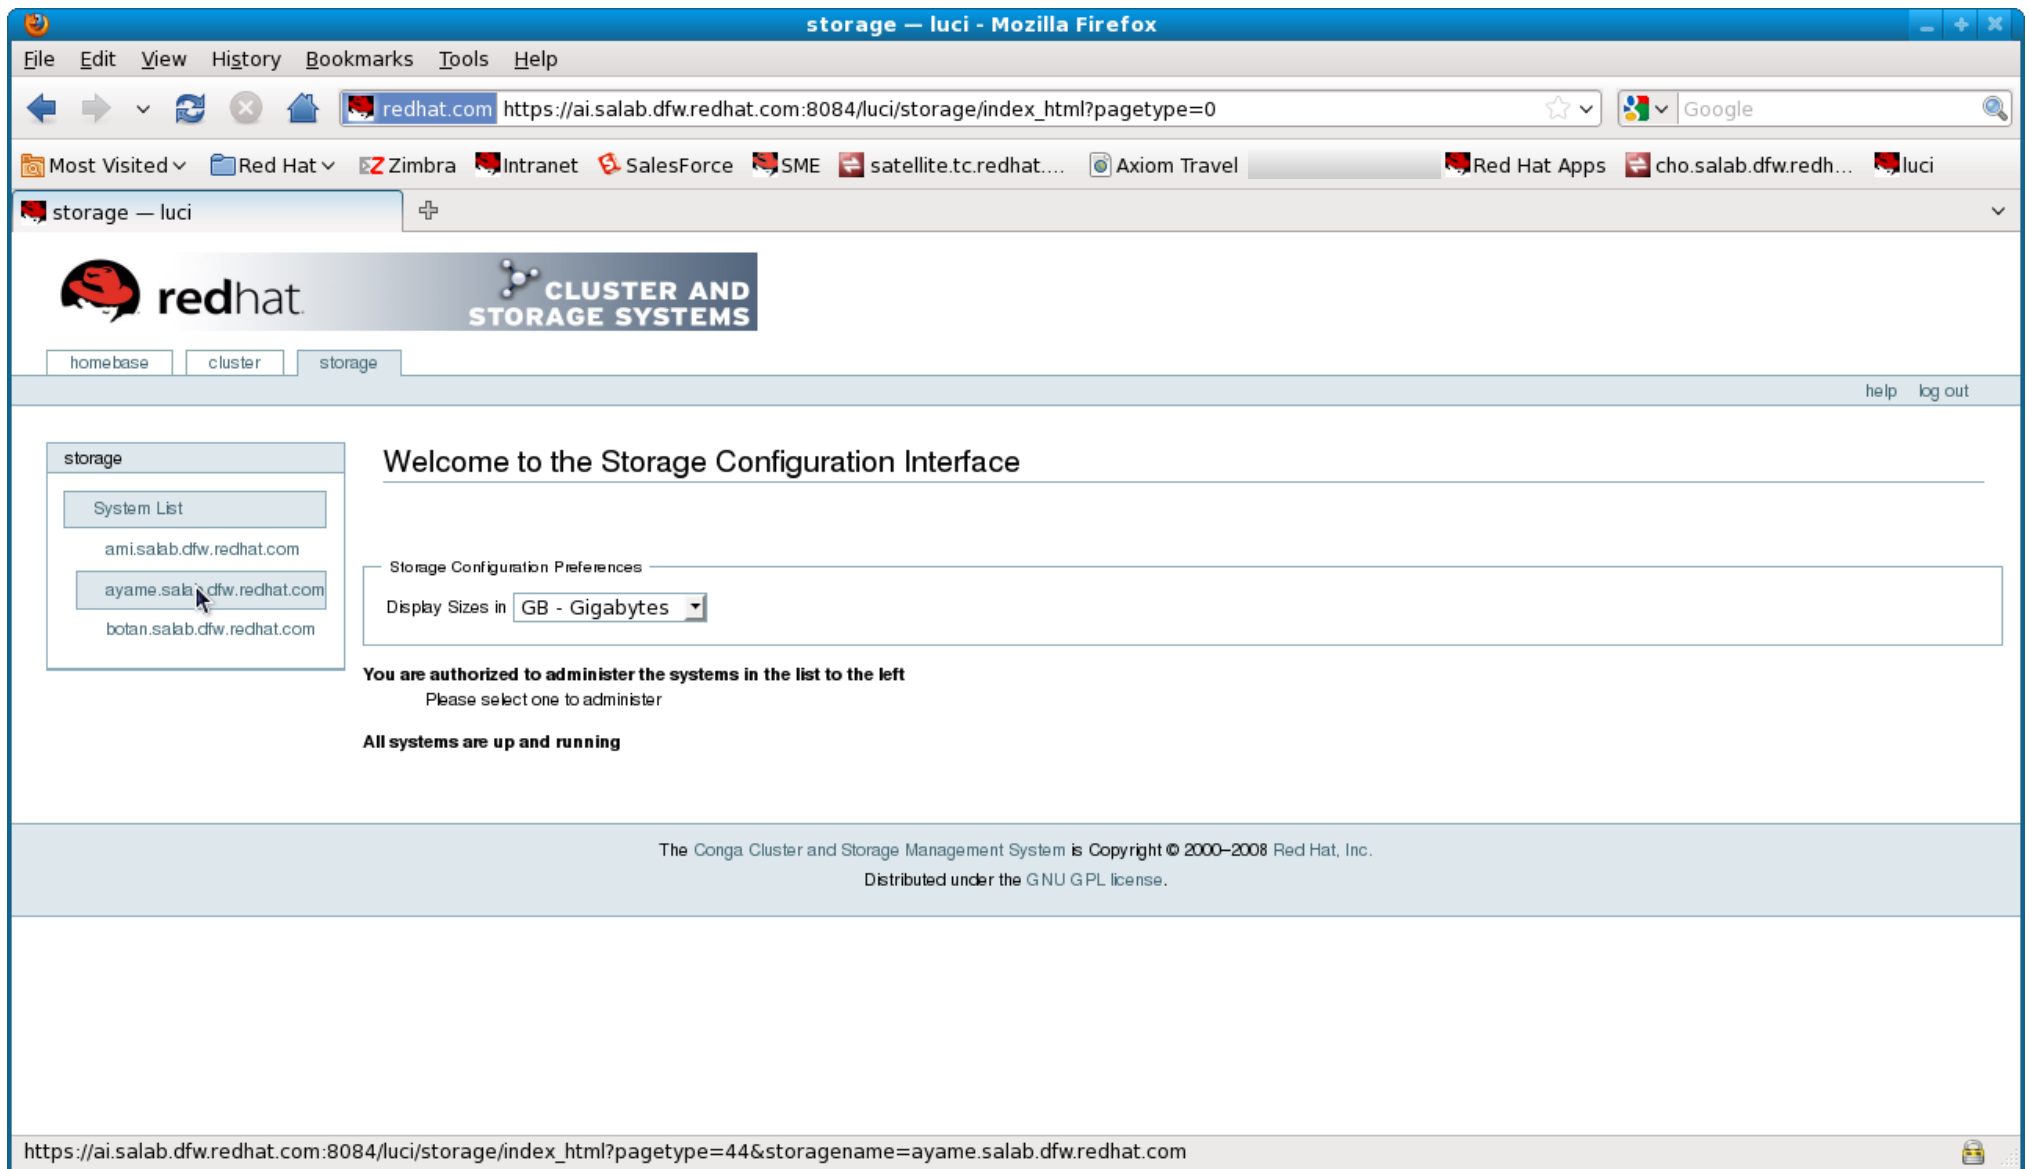
System name

Expand VG

Expand LV

Fill in Mountpoint

Fill in /etc/fstab Mountpoint



storage — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/storage/index_html?pagetype=44&storagename=ayame.salab.df

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

storage — luci

Volume Groups

- VolGroup00
- XenVG
 - Volume Group Properties:
 - Logical Volumes:
 - XenLV 99.98 GB

Logical Volume 'XenLV' - /dev/XenVG/XenLV

| | |
|---------------------|--|
| Logical Volume Name | XenLV |
| Volume Group Name | XenVG |
| Extent Size | 4.0 MB |
| Size | 99.98 GB |
| Mirrored | false |
| Attributes | -wi-a- |
| Clustered | true |
| Snapshots | |
| UUID | XQ2fss-WDvo-KkNj-IRM3-1n2h-Fd5d-zxWSTC |

Content GFS2 - Global FS v.2

| | |
|-----------------------|---------------------|
| Cluster Name | dom0-cluster |
| Unique GFS Name | gfs0 |
| Block Size | 4.0 KB |
| Mountpoint | /var/lib/xen/images |
| /etc/fstab Mountpoint | /var/lib/xen/images |
| Mountable | true |
| Locking Protocol | dlm |
| Clustered | true |

Remove Reset Apply

Done

Configure Storage (alternative method)

Copy and paste the last line of /etc/fstab to the other nodes and then mount -a

```
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]# mount  
/dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/sda1 on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /var/lib/xenstored type tmpfs (rw)  
none on /sys/kernel/config type configfs (rw)  
[root@botan ~]# echo /dev/XenVG/XenLV /var/lib/xen/images gfs2 defaults 0 0 >> /  
etc/fstab  
[root@botan ~]# mount -a  
[root@botan ~]# mount  
/dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/sda1 on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /var/lib/xenstored type tmpfs (rw)  
none on /sys/kernel/config type configfs (rw)  
/dev/mapper/XenVG-XenLV on /var/lib/xen/images type gfs2 (rw,hostdata=jid=2:id=1  
31073:first=0)  
[root@botan ~]#
```

Verify Storage Availability

Something simple like a `df -h` on all nodes is fine.


```
root@ai:~  
File Edit View Terminal Help  
[root@ai ~]#  
  
root@ayame:~  
File Edit View Terminal Help  
[root@ayame ~]# df -h  
Filesystem      Size  Used Avail Use% Mounted on  
/dev/mapper/VolGroup00-LogVol00 901G  2.9G  851G   1% /  
/dev/sda1        99M   48M   46M  52% /boot  
tmpfs            3.7G    0  3.7G   0% /dev/shm  
none            3.7G  104K  3.7G   1% /var/lib/xenstored  
/dev/mapper/XenVG-XenLV 100G  166M  100G   1% /var/lib/xen/images  
[root@ayame ~]#  
  
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# df -h  
Filesystem      Size  Used Avail Use% Mounted on  
/dev/mapper/VolGroup00-LogVol00 901G  39G  815G   5% /  
/dev/sda1        99M   48M   46M  52% /boot  
tmpfs            3.7G    0  3.7G   0% /dev/shm  
none            3.7G  104K  3.7G   1% /var/lib/xenstored  
/dev/mapper/XenVG-XenLV 100G  166M  100G   1% /var/lib/xen/images  
[root@ami ~]#  
  
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]# df -h  
Filesystem      Size  Used Avail Use% Mounted on  
/dev/mapper/VolGroup00-LogVol00 901G  2.9G  851G   1% /  
/dev/sda1        99M   48M   46M  52% /boot  
tmpfs            3.7G    0  3.7G   0% /dev/shm  
none            3.7G  104K  3.7G   1% /var/lib/xenstored  
/dev/mapper/XenVG-XenLV 100G  166M  100G   1% /var/lib/xen/images  
[root@botan ~]#
```

Prepare dom0 machines for live migration

Make sure that time is synced across all of the dom0 machines.

You **must** use something like NTP. Even if you just set up the management server as an NTP server, the dom0 times **must** be synced for live migration to work right!

The image displays four terminal windows arranged in a 2x2 grid, each showing the output of the command `grep -v "^#" /etc/ntp.conf | uniq`. The windows are titled `root@ai:~`, `root@ayame:~`, `root@ami:~`, and `root@botan:~`. Each terminal has a menu bar with `File`, `Edit`, `View`, `Terminal`, and `Help`. The output in each terminal is as follows:

```
root@ai:~# grep -v "^#" /etc/ntp.conf | uniq
restrict default kod nomodify notrap nopeer noquery

restrict 127.0.0.1

driftfile /var/lib/ntp/drift

keys /etc/ntp/keys

server 127.127.1.0
fudge 127.127.1.0 stratum 10
server clock.util.phx2.redhat.com
restrict clock.util.phx2.redhat.com mask 255.255.255.255 nomodify notrap noquery
root@ai ~#
```

```
root@ayame:~# grep -v "^#" /etc/ntp.conf | uniq
restrict default kod nomodify notrap nopeer noquery

restrict 127.0.0.1

driftfile /var/lib/ntp/drift

keys /etc/ntp/keys

restrict clock.util.phx2.redhat.com mask 255.255.255.255 nomodify notrap noquery
server 127.127.1.0
fudge 127.127.1.0 stratum 10
server clock.util.phx2.redhat.com
root@ayame ~#
```

```
root@ami:~# grep -v "^#" /etc/ntp.conf | uniq
restrict default kod nomodify notrap nopeer noquery

restrict 127.0.0.1

driftfile /var/lib/ntp/drift

keys /etc/ntp/keys

restrict clock.util.phx2.redhat.com mask 255.255.255.255 nomodify notrap noquery
server 127.127.1.0
fudge 127.127.1.0 stratum 10
server clock.util.phx2.redhat.com
root@ami ~#
```

```
root@botan:~# grep -v "^#" /etc/ntp.conf | uniq
restrict default kod nomodify notrap nopeer noquery

restrict 127.0.0.1

driftfile /var/lib/ntp/drift

keys /etc/ntp/keys

restrict clock.util.phx2.redhat.com mask 255.255.255.255 nomodify notrap noquery
server 127.127.1.0
fudge 127.127.1.0 stratum 10
server clock.util.phx2.redhat.com
root@botan ~#
```

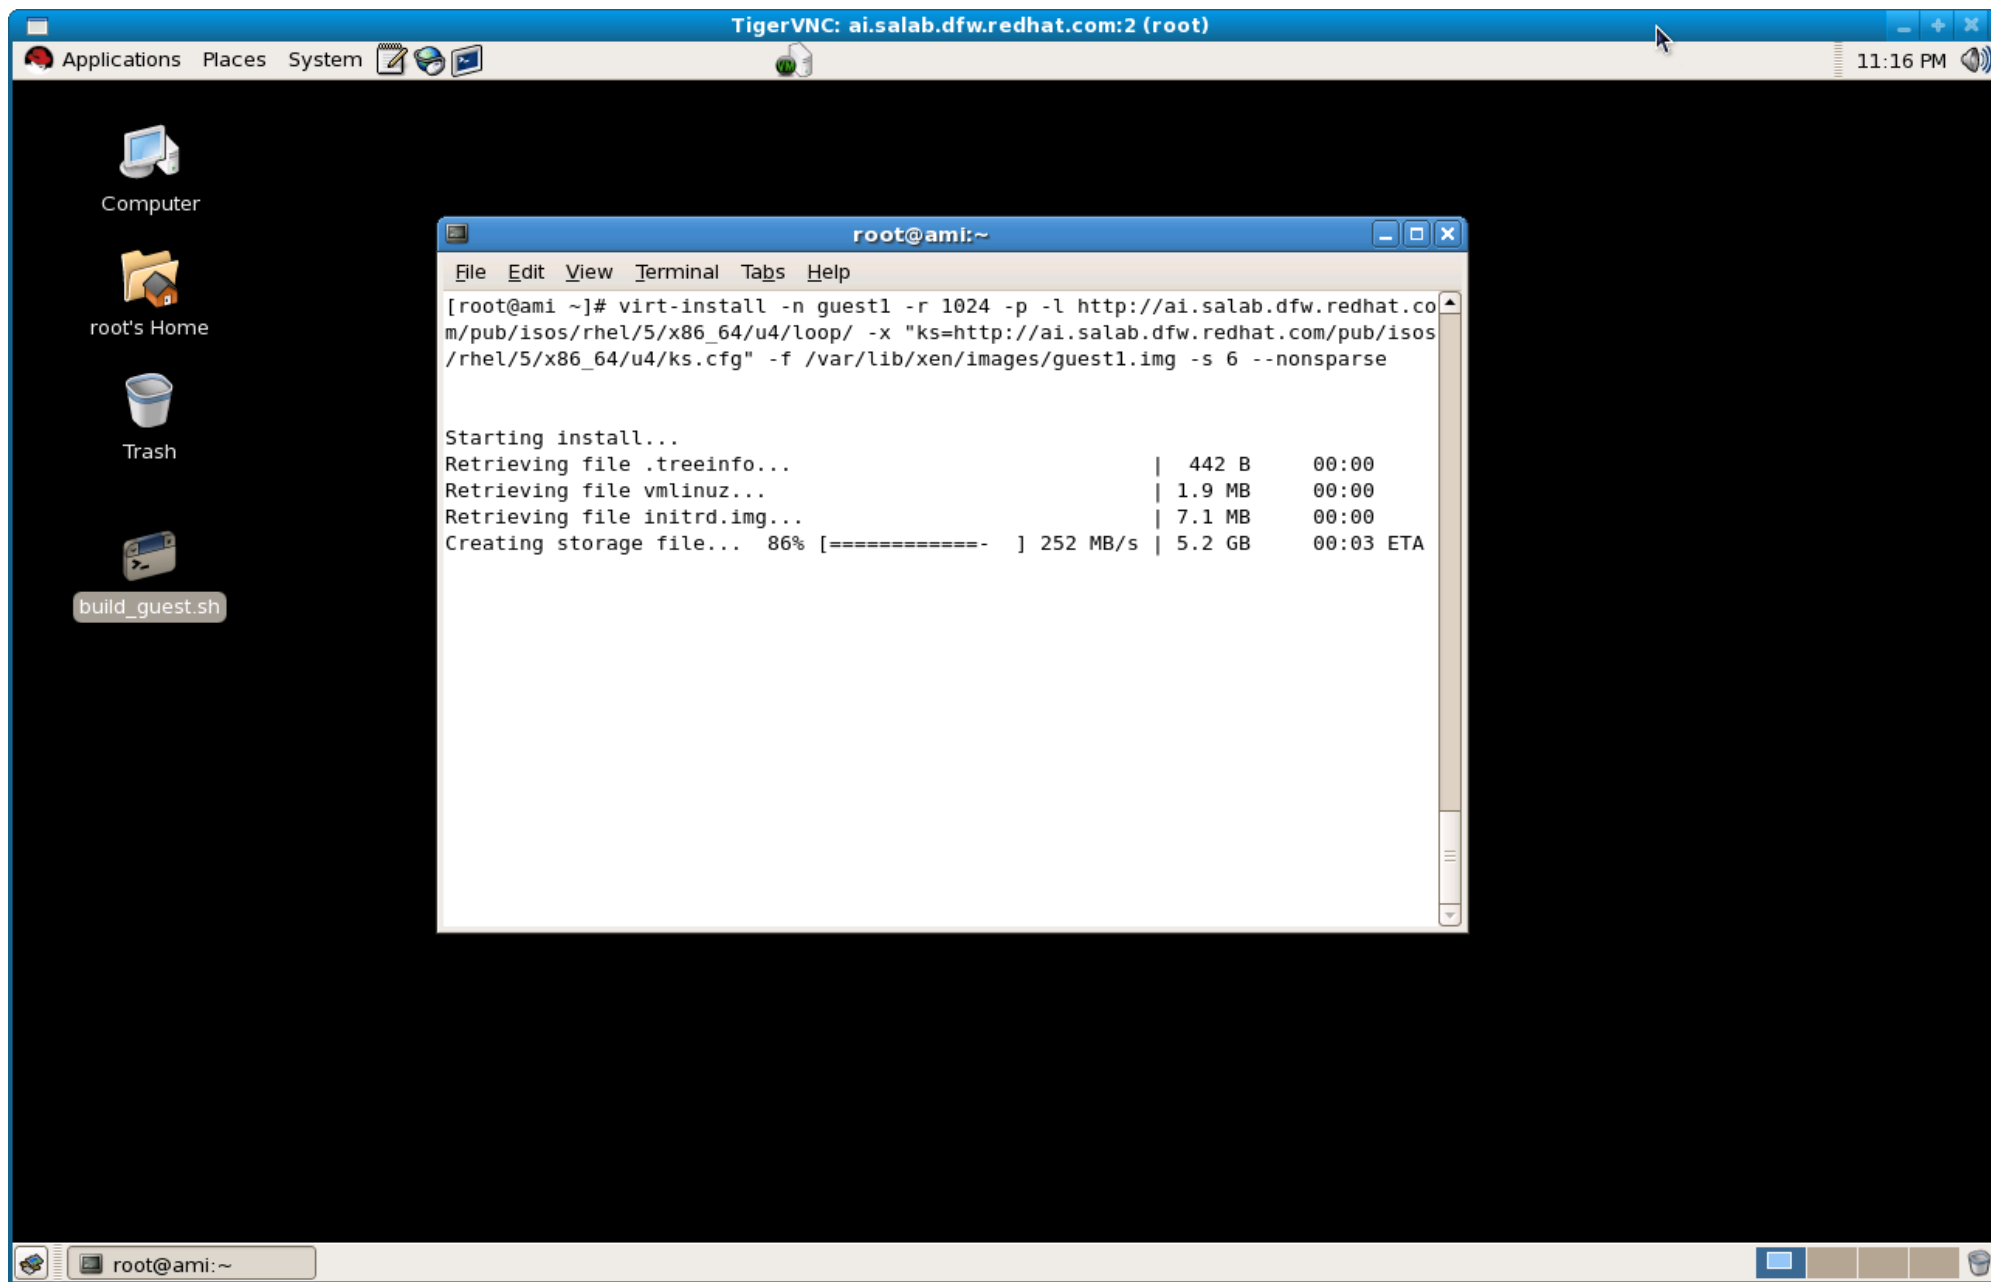
Install Virtual Machines

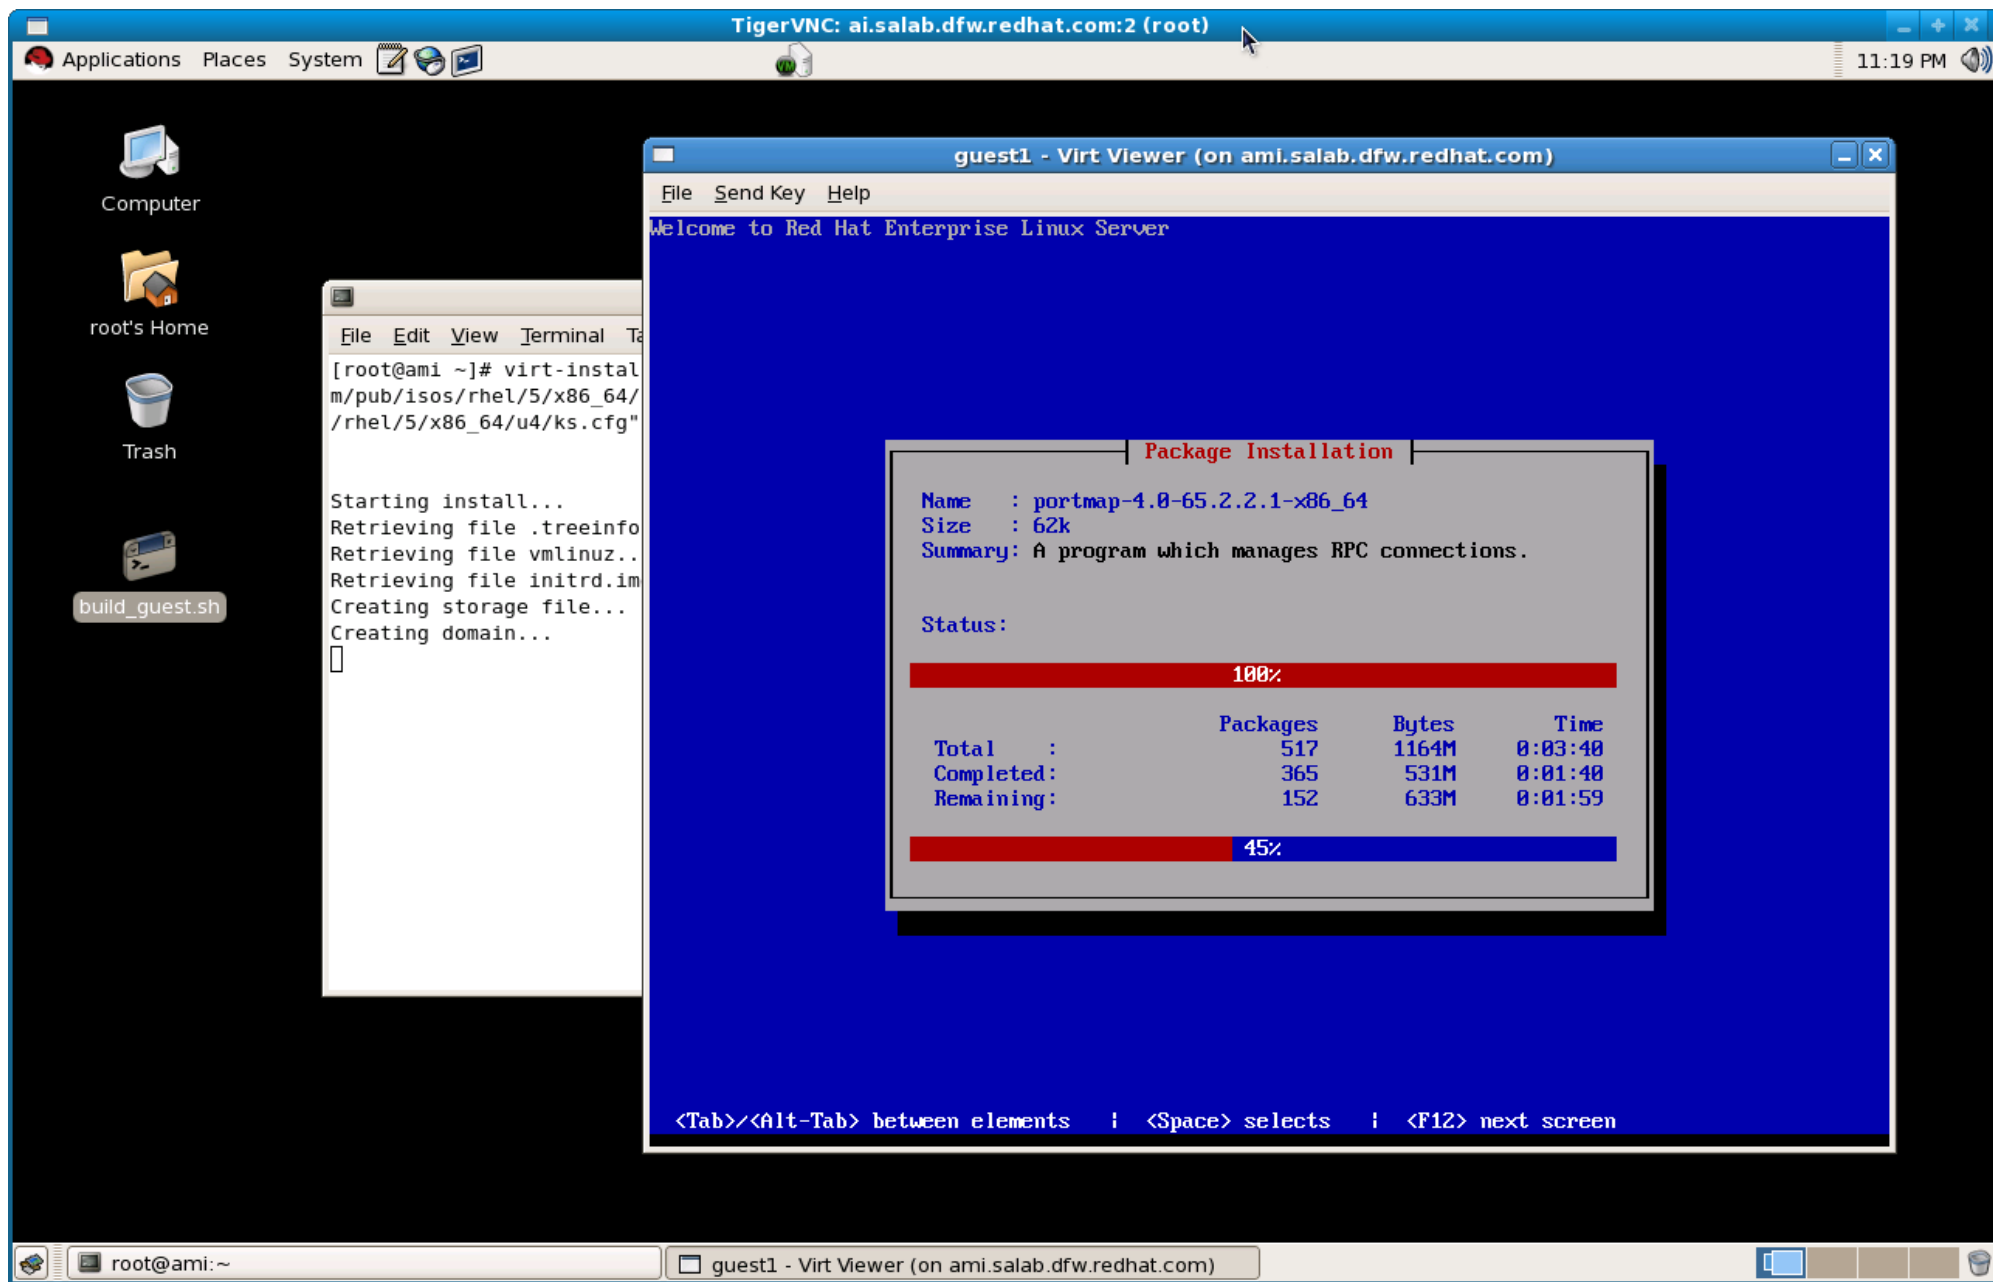
Use virt-manager or virt-install

The following is a single command, separated here for easier viewing:

```
virt-install -n guest1 -r 1024 -p \  
-l http://ai.salab.dfw.redhat.com/pub/isos/rhel/5/x86_64/u4/loop/ \  
-x "ks=http://ai.salab.dfw.redhat.com/pub/isos/rhel/5/x86_64/u4/ks.cfg" \  
-f /var/lib/xen/images/guest1.img -s 6 --nonsparse
```

^^^^^^^^^^





Prepare all dom0 machines for live migration

In the `/etc/xend-config.sxp` file set the following:

`(xend-relocation-server yes)`

`(xend-relocation-port 8002)`

`(xend-relocation-address "")`

`(xend-relocation-hosts-allow "")`

Restart the xend service

Test live migration

```
virsh migrate guest1 \ xen  
+ssh://root@ayame.salab.dfw.redhat.com
```

```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# grep xend-relocation /etc/xen/xend-config.sxp | grep -v "^#"  
(xend-relocation-server yes)  
(xend-relocation-port 8002)  
(xend-relocation-address '')  
(xend-relocation-hosts-allow '')  
[root@ami ~]# service xend restart  
restart xend: [ OK ]  
[root@ami ~]# virsh migrate guest1 xen+ssh://root@ayame.salab.dfw.redhat.com  
root@ayame.salab.dfw.redhat.com's password:  
  
[root@ami ~]# virsh --connect xen+ssh://root@ayame.salab.dfw.redhat.com  
root@ayame.salab.dfw.redhat.com's password:  
Welcome to virsh, the virtualization interactive terminal.  
  
Type: 'help' for help with commands  
      'quit' to quit  
  
virsh # list  
Id Name State  
-----  
0 Domain-0 running  
4 guest1 idle  
  
virsh #
```


Prepare dom0 machines for live migration

Stop all the guests

`chkconfig xendomains off`

`service xendomains stop`



A terminal window titled "root@ami:~" with a menu bar (File, Edit, View, Terminal, Help). The terminal shows the following commands and output:

```
[root@ami ~]# virsh list
Id Name                               State
-----
 0 Domain-0                           running

[root@ami ~]# chkconfig xendomains off
[root@ami ~]# service xendomains stop
Shutting down Xen domains:[done]      [ OK ]
[root@ami ~]#
```

A mouse cursor is visible over the prompt line.

Prepare dom0 machines for live migration

Normally you would move the config files etc /etc/xen to shared storage (/var/lib/xen/images on each dom0), but there is a bug in the beta version of RHEL 5.4, BZ 519786.

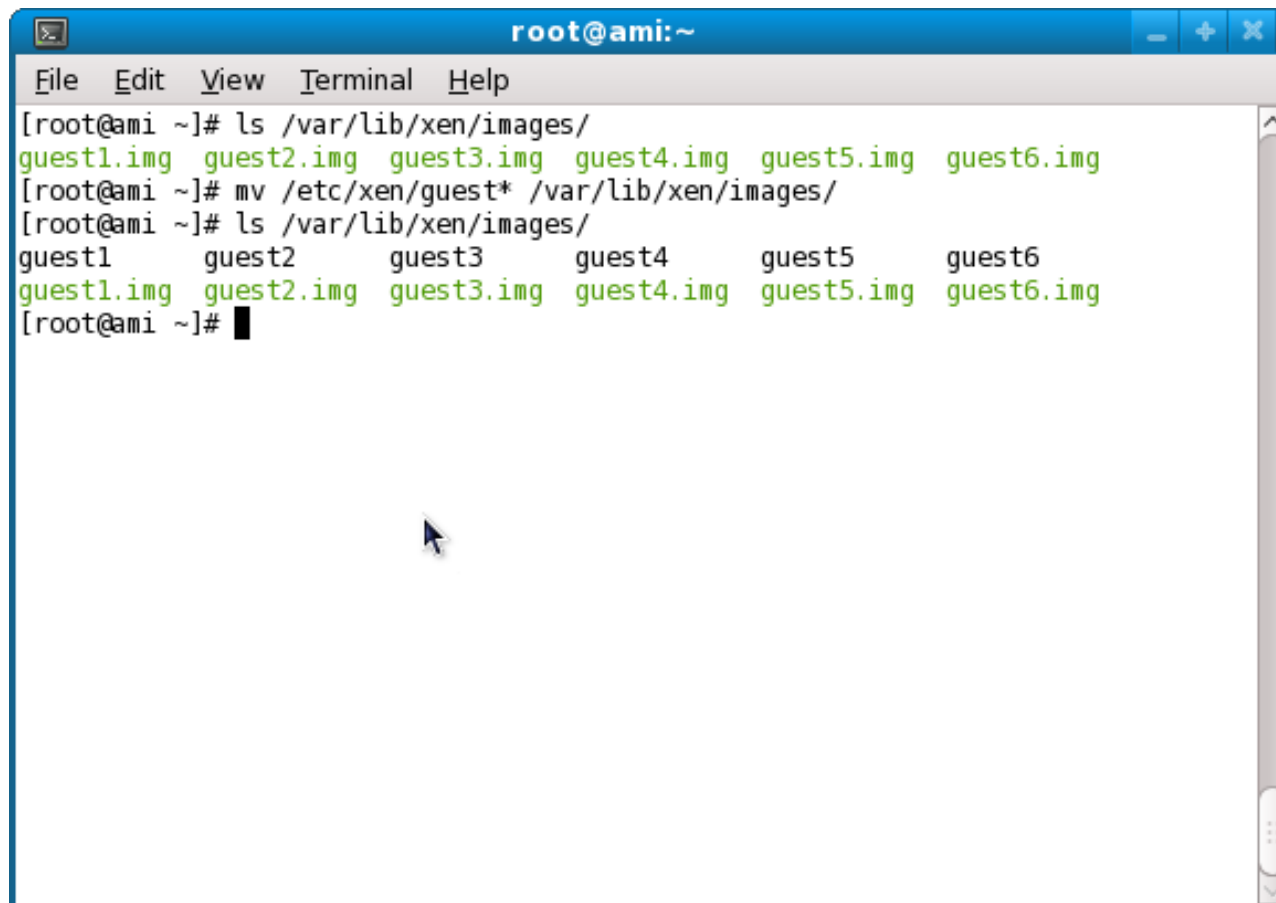
Either:

- use the vm.sh from the BZ

- copy all of the domU config files to /etc/xen on all of the nodes

- symlink the configs from /etc/xen/[guest] to /var/lib/xen/images/[guest] on all the nodes.

For this presentation, I used the vm.sh from the BZ and moved all the configs to shared storage

A terminal window titled 'root@ami:~' with a menu bar (File, Edit, View, Terminal, Help). The terminal shows the following commands and output:

```
[root@ami ~]# ls /var/lib/xen/images/  
guest1.img guest2.img guest3.img guest4.img guest5.img guest6.img  
[root@ami ~]# mv /etc/xen/guest* /var/lib/xen/images/  
[root@ami ~]# ls /var/lib/xen/images/  
guest1      guest2      guest3      guest4      guest5      guest6  
guest1.img  guest2.img  guest3.img  guest4.img  guest5.img  guest6.img  
[root@ami ~]#
```

A mouse cursor is visible in the center of the terminal window.

Configure Failover Domains

Cluster tab

Cluster Name

Failover Domains

Add Failover Domain

FD Name

Check “Prioritized”

Check “Restrict failover to this domain's members”

Do not check “Do not fail back services...”

Assign members and priorities

Submit

Luci — cluster — failover domains — Add a failover domain - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=41&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — failover domains...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
 - Add a Failover Domain
 - Configure a Failover Domain
- Shared Fence Devices

dom0-cluster

Add a Failover Domain

Failover Domain Name prefer_ami

Prioritized ☒

Restrict failover to this domain's members ☒

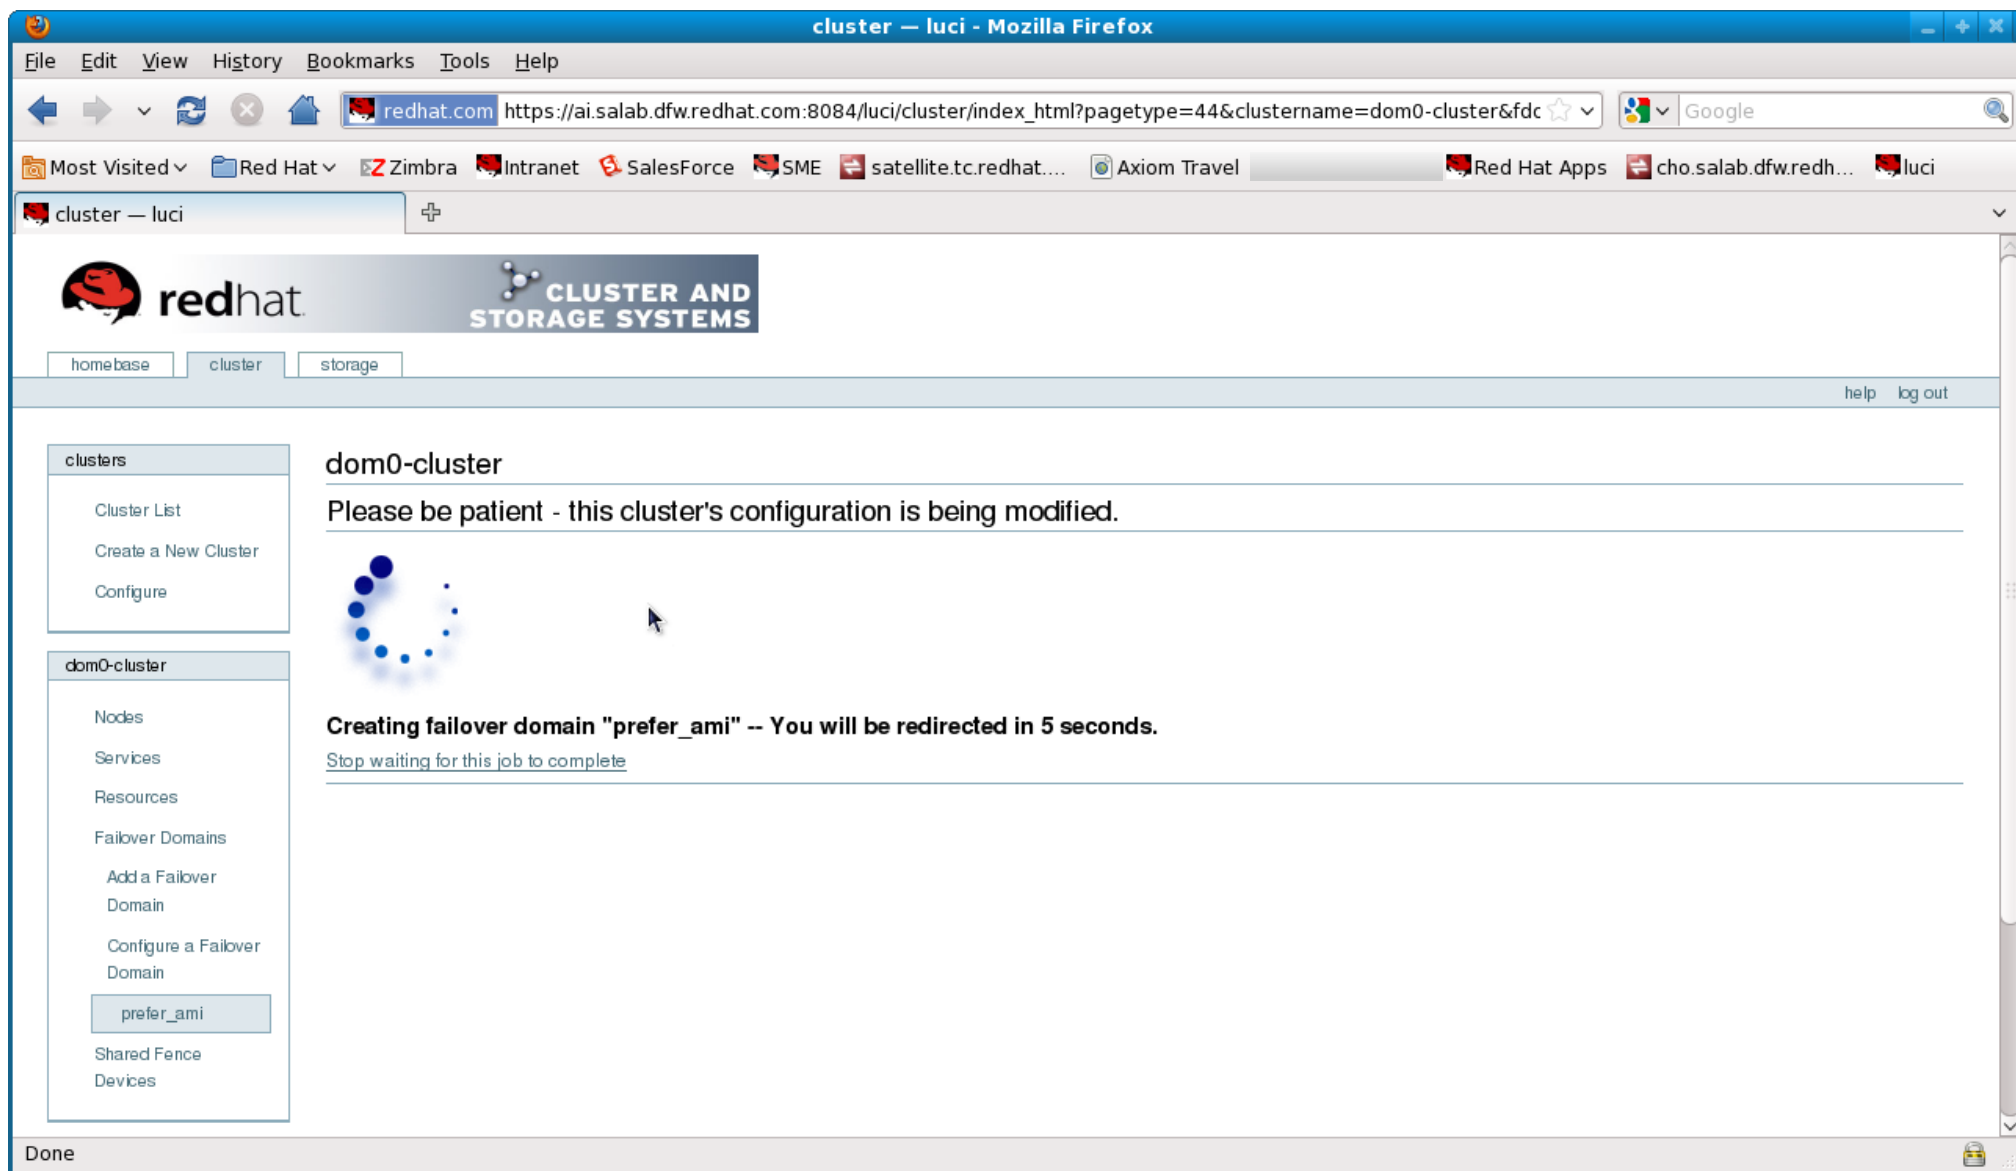
Do not fail back services in this domain ☐

Failover domain membership

| Node | Member | Priority |
|----------------------------|-------------------------------------|----------|
| ami.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | 1 |
| ayame.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | 10 |
| botan.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | 20 |

Submit

Done



Configure Failover Domains

Create a domain for each dom0.

prefer_ami

ami 1, ayame 10, botan 20

prefer_ayame

ayame 1, botan 10, ami 20

prefer_botan

botan 1, ami 10, ayame 20

Luci — cluster — failover domains - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=40&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — failover domains

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
 - Add a Failover Domain
 - Configure a Failover Domain
 - Shared Fence Devices

dom0-cluster

Failover Domain Name: [prefer_ami](#)

This Failover Domain is Ordered

This Failover Domain is Restricted

Members of this Failover Domain

- [ami.salab.dfw.redhat.com](#)
- [ayame.salab.dfw.redhat.com](#)
- [botan.salab.dfw.redhat.com](#)

Services employing this Failover Domain:

- No Services Defined

Failover Domain Name: [prefer_ayame](#)

This Failover Domain is Ordered

This Failover Domain is Restricted

Members of this Failover Domain

- [ami.salab.dfw.redhat.com](#)

Done

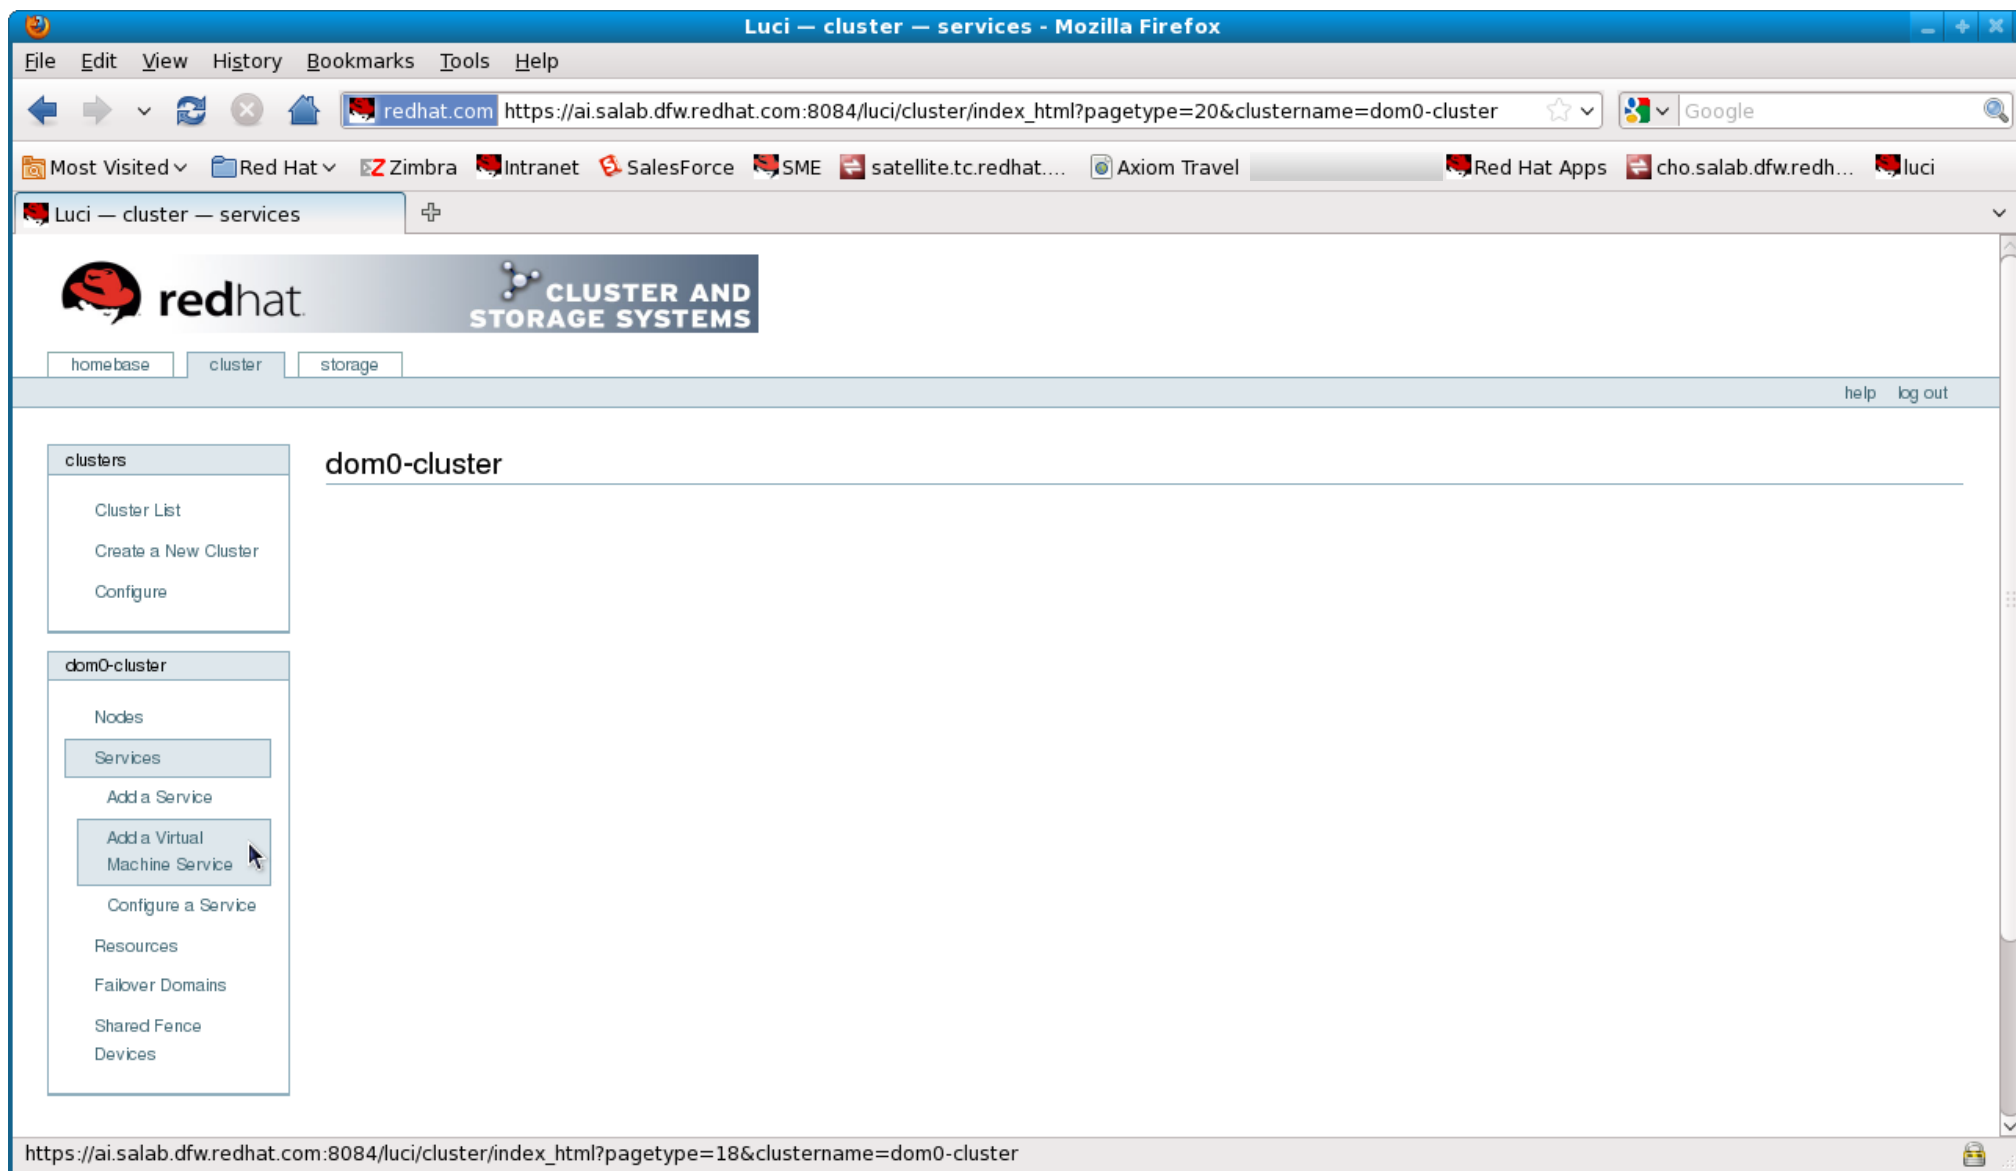
Define Virtual Machine Services

Cluster Tab

Cluster Name

Services

Add Virtual Machine Service



cluster — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?clustername=dom0-cluster&servicename=gue

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

cluster — luci

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
- Services
 - Add a Service
 - Add a Virtual Machine Service
 - Configure a Service
 - guest1
 - guest2
 - guest3
 - guest4
 - guest5

dom0-cluster

Properties for guest1

| | |
|---|-------------------------------------|
| Virtual machine name | guest1 |
| Path to VM configuration files | /var/lib/xen/images |
| VM Migration Mapping | |
| Migration type | Live |
| Hypervisor | Xen |
| Automatically start this service | <input checked="" type="checkbox"/> |
| Enable NFS lock workarounds | <input type="checkbox"/> |
| Run exclusive | <input type="checkbox"/> |
| Failover Domain | prefer_ami |
| Recovery policy | Restart |
| Maximum number of restart failures before relocating | 1 |
| Length of time in seconds after which to forget a restart | 600 |

Update Virtual Machine Service Delete Virtual Machine Service

Done

Define Virtual Machine Services

Repeat for each VM Service

In this example, two VM Services are assigned to each failover domain - this equates to two machines starting on ai, two on ayame and two on botan

Luci — cluster — failover domains - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=40&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — failover do... clurgmgrd: <warning> #68: ... Bug 236276 - rgmanager fail...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
- Services
- Resources
- Failover Domains
 - Add a Failover Domain
 - Configure a Failover Domain
 - Shared Fence Devices

dom0-cluster

Failover Domain Name: [prefer_ami](#)

This Failover Domain is Ordered

This Failover Domain is Restricted

Members of this Failover Domain

- [ami.salab.dfw.redhat.com](#)
- [ayame.salab.dfw.redhat.com](#)
- [botan.salab.dfw.redhat.com](#)

Services employing this Failover Domain:

- guest1 Running on node [ami.salab.dfw.redhat.com](#)
- guest2 Running on node [ami.salab.dfw.redhat.com](#)

Failover Domain Name: [prefer_ayame](#)

This Failover Domain is Ordered

This Failover Domain is Restricted

Members of this Failover Domain

Done

Test virtual machine management services

Cluster Tab

Cluster Name

Services

Choose an action for one of the machines from the drop
down menu

Go and confirm

Luci — cluster — services - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=20&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
- Services
- Add a Service
- Add a Virtual Machine Service
- Configure a Service
- Resources
- Failover Domains
- Shared Fence
- Devices

dom0-cluster

Service Name guest1 (virtual machine service)

Status: Running on ami.salab.dfw.redhat.com
Autostart is enabled for this service

Failover Domain Association: prefer_ami

Relocate this service to ayame.salab.dfw.redhat.com Go

Service Name guest2 (virtual machine service)

Status: Running on ami.salab.dfw.redhat.com
Autostart is enabled for this service

Failover Domain Association: prefer_ami

Choose a Task... Go

Service Name guest3 (virtual machine service)

Status: Running on ayame.salab.dfw.redhat.com
Autostart is enabled for this service

Failover Domain Association: prefer_ayame

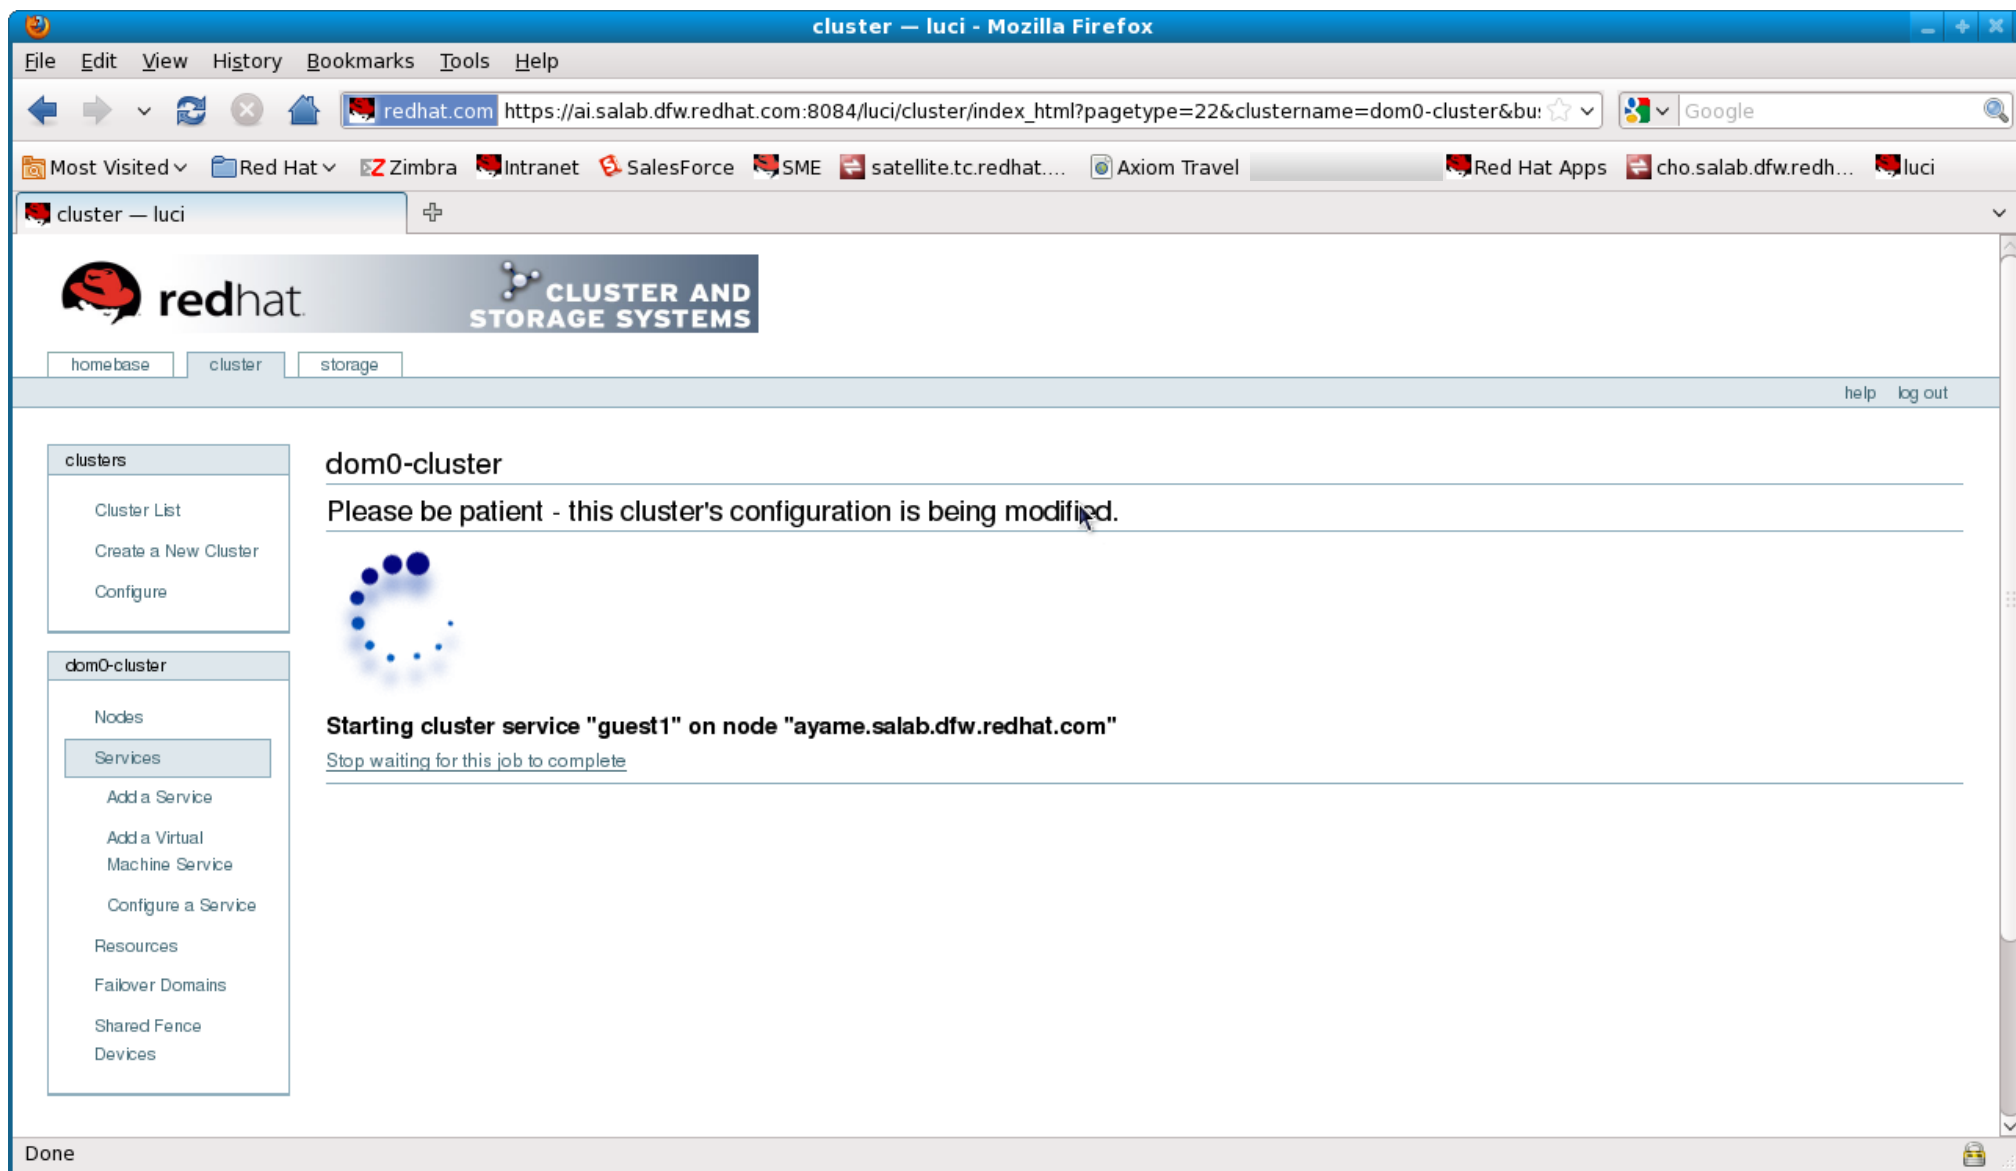
Choose a Task... Go

Service Name guest4 (virtual machine service)

Status: Running on ayame.salab.dfw.redhat.com
Autostart is enabled for this service

Choose a Task... Go

Done



Luci — cluster — services - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=22&clustername=dom0-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

- Nodes
- Services
- Add a Service
- Add a Virtual Machine Service
- Configure a Service
- Resources
- Failover Domains
- Shared Fence
- Devices

dom0-cluster

Service Name guest2 (virtual machine service)

Status: Running on ami.salab.dfw.redhat.com
Autostart is enabled for this service

Failover Domain Association: prefer_ami

Service Name guest1 (virtual machine service)

Status: Running on ayame.salab.dfw.redhat.com
Autostart is enabled for this service

Failover Domain Association: prefer_ami

Service Name guest3 (virtual machine service)

Status: Running on ayame.salab.dfw.redhat.com
Autostart is enabled for this service

Failover Domain Association: prefer_ayame

Service Name guest4 (virtual machine service)

Status: Running on ayame.salab.dfw.redhat.com
Autostart is enabled for this service

Done

```
root@ayame:~  
File Edit View Terminal Help  
[root@ayame ~]# virsh list  
Id Name State  
-----  
0 Domain-0 running  
5 guest4 idle  
6 guest3 idle  
8 guest1 idle  
  
[root@ayame ~]# clustat  
Cluster Status for dom0-cluster @ Thu Aug 27 17:10:35 2009  
Member Status: Quorate  
  
Member Name ID Status  
-----  
ami.salab.dfw.redhat.com 1 Online, rgmanager  
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com 3 Online, rgmanager  
  
Service Name Owner (Last) State  
-----  
vm:guest1 ayame.salab.dfw.redhat.com started  
vm:guest2 ami.salab.dfw.redhat.com started  
vm:guest3 ayame.salab.dfw.redhat.com started  
vm:guest4 ayame.salab.dfw.redhat.com started  
vm:guest5 botan.salab.dfw.redhat.com started  
vm:guest6 botan.salab.dfw.redhat.com started  
[root@ayame ~]#
```

Test virtual machine failover services

Log into a domU machine and crash it

```
echo c > /proc/sysrq-trigger
```

Watch clustat on one of the dom0 machines

In the following slides, the domU system is crashed and the cluster restarts it on ami.

```
root@sayoko:~  
File Edit View Terminal Help  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 10:56:25 2009  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
█  
  
root@ayame:~  
File Edit View Terminal Help  
Every 2.0s: clustat Thu Aug 27 17:26:40 2009  
Cluster Status for dom0-cluster @ Thu Aug 27 17:26:40 2009  
Member Status: Quorate  
  
Member Name ID Status  
-----  
ami.salab.dfw.redhat.com 1 Online, rgmanager  
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com 3 Online, rgmanager  
  
Service Name Owner (Last) State  
-----  
vm:guest1 ami.salab.dfw.redhat.com started  
vm:guest2 ami.salab.dfw.redhat.com started  
vm:guest3 ayame.salab.dfw.redhat.com started  
vm:guest4 ayame.salab.dfw.redhat.com started  
vm:guest5 botan.salab.dfw.redhat.com started  
vm:guest6 botan.salab.dfw.redhat.com started  
  
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# virsh console guest1  
Connected to domain guest1  
Escape character is ^]  
  
Red Hat Enterprise Linux Server release 5.4 (Tikanga)  
Kernel 2.6.18-164.el5xen on an x86_64  
  
sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump  
Kexec: Warning: crash image not loaded  
Kernel panic - not syncing: SysRq-triggered panic!  
█  
  
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]# █
```

```
root@sayoko:~  
File Edit View Terminal Help  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 10:56:25 2009  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
[root@sayoko ~]#  
  
root@ayame:~  
File Edit View Terminal Help  
Every 2.0s: clustat Thu Aug 27 17:26:58 2009  
Cluster Status for dom0-cluster @ Thu Aug 27 17:26:58 2009  
Member Status: Quorate  


| Member Name                | ID | Status                   |
|----------------------------|----|--------------------------|
| ami.salab.dfw.redhat.com   | 1  | Online, rgmanager        |
| ayame.salab.dfw.redhat.com | 2  | Online, Local, rgmanager |
| botan.salab.dfw.redhat.com | 3  | Online, rgmanager        |


| Service Name | Owner (Last)                 | State       |
|--------------|------------------------------|-------------|
| vm:guest1    | (ayame.salab.dfw.redhat.com) | recoverable |
| vm:guest2    | ami.salab.dfw.redhat.com     | started     |
| vm:guest3    | ayame.salab.dfw.redhat.com   | started     |
| vm:guest4    | ayame.salab.dfw.redhat.com   | started     |
| vm:guest5    | botan.salab.dfw.redhat.com   | started     |
| vm:guest6    | botan.salab.dfw.redhat.com   | started     |

  
  
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# virsh console guest1  
Connected to domain guest1  
Escape character is ^]  
  
Red Hat Enterprise Linux Server release 5.4 (Tikanga)  
Kernel 2.6.18-164.el5xen on an x86_64  
  
sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump  
Kexec: Warning: crash image not loaded  
Kernel panic - not syncing: SysRq-triggered panic!  
  
[root@ami ~]#  
  
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]#
```

```
root@sayoko:~  
File Edit View Terminal Help  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 10:56:25 2009  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
[  
root@ayame:~  
File Edit View Terminal Help  
Every 2.0s: clustat Thu Aug 27 17:27:02 2009  
Cluster Status for dom0-cluster @ Thu Aug 27 17:27:02 2009  
Member Status: Quorate  
  
Member Name ID Status  
-----  
ami.salab.dfw.redhat.com 1 Online, rgmanager  
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com 3 Online, rgmanager  
  
Service Name Owner (Last) State  
-----  
vm:guest1 none recovering  
vm:guest2 ami.salab.dfw.redhat.com started  
vm:guest3 ayame.salab.dfw.redhat.com started  
vm:guest4 ayame.salab.dfw.redhat.com started  
vm:guest5 botan.salab.dfw.redhat.com started  
vm:guest6 botan.salab.dfw.redhat.com started  
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# virsh console guest1  
Connected to domain guest1  
Escape character is ^]  
  
Red Hat Enterprise Linux Server release 5.4 (Tikanga)  
Kernel 2.6.18-164.el5xen on an x86_64  
  
sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump  
Kexec: Warning: crash image not loaded  
Kernel panic - not syncing: SysRq-triggered panic!  
  
[root@ami ~]#  
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]#
```

```
root@sayoko:~
File Edit View Terminal Help
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 10:56:25 2009
[root@sayoko ~]# echo c > /proc/sysrq-trigger
[]

root@ayame:~
File Edit View Terminal Help
Every 2.0s: clustat Thu Aug 27 17:27:04 2009
Cluster Status for dom0-cluster @ Thu Aug 27 17:27:04 2009
Member Status: Quorate

Member Name              ID   Status
-----
ami.salab.dfw.redhat.com 1   Online, rgmanager
ayame.salab.dfw.redhat.com 2   Online, Local, rgmanager
botan.salab.dfw.redhat.com 3   Online, rgmanager

Service Name              Owner (Last)              State
-----
vm:guest1                 ami.salab.dfw.redhat.com  started
vm:guest2                 ami.salab.dfw.redhat.com  started
vm:guest3                 ayame.salab.dfw.redhat.com started
vm:guest4                 ayame.salab.dfw.redhat.com started
vm:guest5                 botan.salab.dfw.redhat.com started
vm:guest6                 botan.salab.dfw.redhat.com started

root@ami:~
File Edit View Terminal Help
[root@ami ~]# virsh console guest1
Connected to domain guest1
Escape character is ^]

Red Hat Enterprise Linux Server release 5.4 (Tikanga)
Kernel 2.6.18-164.el5xen on an x86_64

sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump
Kexec: Warning: crash image not loaded
Kernel panic - not syncing: SysRq-triggered panic!

[root@ami ~]# []

root@botan:~
File Edit View Terminal Help
[root@botan ~]# []
```


Test virtual machine failover services

The same guest is crashed a second time and the cluster migrates it to ayami based on the failover domain rules we set up earlier.

```
root@sayoko:~
File Edit View Terminal Help
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 10:56:25 2009
[root@sayoko ~]# echo c > /proc/sysrq-trigger
Read from remote host sayoko.salab.dfw.redhat.com: Connection reset by peer
Connection to sayoko.salab.dfw.redhat.com closed.
[root@ai ~]#
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 17:23:27 2009 from ai.salab.dfw.redhat.com
[root@sayoko ~]# echo c > /proc/sysrq-trigger

```

```
root@ayame:~
File Edit View Terminal Help
Every 2.0s: clustat Thu Aug 27 17:27:54 2009

Cluster Status for dom0-cluster @ Thu Aug 27 17:27:54 2009
Member Status: Quorate

Member Name              ID  Status
-----
ami.salab.dfw.redhat.com  1  Online, rgmanager
ayame.salab.dfw.redhat.com 2  Online, Local, rgmanager
botan.salab.dfw.redhat.com 3  Online, rgmanager

Service Name              Owner (Last)              State
-----
vm:guest1                 ami.salab.dfw.redhat.com  started
vm:guest2                 ami.salab.dfw.redhat.com  started
vm:guest3                 ayame.salab.dfw.redhat.com started
vm:guest4                 ayame.salab.dfw.redhat.com started
vm:guest5                 botan.salab.dfw.redhat.com started
vm:guest6                 botan.salab.dfw.redhat.com started

```

```
root@ami:~
File Edit View Terminal Help
device-mapper: uevent: version 1.0.3
device-mapper: ioctl: 4.11.5-ioclt (2007-12-12) initialised: dm-devel@redhat.com
device-mapper: dm-raid45: initialized v0.2594l
EXT3-fs: INFO: recovery required on readonly filesystem.
EXT3-fs: write access will be enabled during recovery.
kjournald starting. Commit interval 5 seconds
EXT3-fs: recovery complete.
EXT3-fs: mounted filesystem with ordered data mode.
SELinux: Disabled at runtime.
type=1404 audit(1251412033.542:2): selinux=0 audit=4294967295 ses=4294967295
Bridge firewalling registered
ip_tables: (C) 2000-2006 Netfilter Core Team
Netfilter messages via NETLINK v0.30.
ip_conntrack version 2.4 (8192 buckets, 65536 max) - 304 bytes per conntrack

Red Hat Enterprise Linux Server release 5.4 (Tikanga)
Kernel 2.6.18-164.el5xen on an x86_64

sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump
Kexec: Warning: crash image not loaded
Kernel panic - not syncing: SysRq-triggered panic!

```

```
root@botan:~
File Edit View Terminal Help
[root@botan ~]#

```

```
root@sayoko:~
File Edit View Terminal Help
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 10:56:25 2009
[root@sayoko ~]# echo c > /proc/sysrq-trigger
Read from remote host sayoko.salab.dfw.redhat.com: Connection reset by peer
Connection to sayoko.salab.dfw.redhat.com closed.
[root@ai ~]#
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 17:23:27 2009 from ai.salab.dfw.redhat.com
[root@sayoko ~]# echo c > /proc/sysrq-trigger

```

```
root@ayame:~
File Edit View Terminal Help
Every 2.0s: clustat Thu Aug 27 17:28:08 2009

Cluster Status for dom0-cluster @ Thu Aug 27 17:28:08 2009
Member Status: Quorate

Member Name              ID   Status
-----
ami.salab.dfw.redhat.com  1   Online, rgmanager
ayame.salab.dfw.redhat.com 2   Online, Local, rgmanager
botan.salab.dfw.redhat.com 3   Online, rgmanager

Service Name              Owner (Last)              State
-----
vm:guest1                 (ami.salab.dfw.redhat.com) recoverable
vm:guest2                 ami.salab.dfw.redhat.com  started
vm:guest3                 ayame.salab.dfw.redhat.com started
vm:guest4                 ayame.salab.dfw.redhat.com started
vm:guest5                 botan.salab.dfw.redhat.com started
vm:guest6                 botan.salab.dfw.redhat.com started

```

```
root@ami:~
File Edit View Terminal Help
device-mapper: uevent: version 1.0.3
device-mapper: ioctl: 4.11.5-ioclt (2007-12-12) initialised: dm-devel@redhat.com
device-mapper: dm-raid45: initialized v0.2594l
EXT3-fs: INFO: recovery required on readonly filesystem.
EXT3-fs: write access will be enabled during recovery.
kjournald starting. Commit interval 5 seconds
EXT3-fs: recovery complete.
EXT3-fs: mounted filesystem with ordered data mode.
SELinux: Disabled at runtime.
type=1404 audit(1251412033.542:2): selinux=0 auid=4294967295 ses=4294967295
Bridge firewalling registered
ip_tables: (C) 2000-2006 Netfilter Core Team
Netfilter messages via NETLINK v0.30.
ip_conntrack version 2.4 (8192 buckets, 65536 max) - 304 bytes per conntrack

Red Hat Enterprise Linux Server release 5.4 (Tikanga)
Kernel 2.6.18-164.el5xen on an x86_64

sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump
Kexec: Warning: crash image not loaded
Kernel panic - not syncing: SysRq-triggered panic!

[root@ami ~]#

```

```
root@botan:~
File Edit View Terminal Help
[root@botan ~]#

```

```
root@sayoko:~  
File Edit View Terminal Help  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 10:56:25 2009  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
Read from remote host sayoko.salab.dfw.redhat.com: Connection reset by peer  
Connection to sayoko.salab.dfw.redhat.com closed.  
[root@ai ~]#  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 17:23:27 2009 from ai.salab.dfw.redhat.com  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
[  
root@ayame:~  
File Edit View Terminal Help  
Every 2.0s: clustat Thu Aug 27 17:28:12 2009  
Cluster Status for dom0-cluster @ Thu Aug 27 17:28:13 2009  
Member Status: Quorate  
  
Member Name ID Status  
-----  
ami.salab.dfw.redhat.com 1 Online, rgmanager  
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com 3 Online, rgmanager  
  
Service Name Owner (Last) State  
-----  
vm:guest1 none recovering  
vm:guest2 ami.salab.dfw.redhat.com started  
vm:guest3 ayame.salab.dfw.redhat.com started  
vm:guest4 ayame.salab.dfw.redhat.com started  
vm:guest5 botan.salab.dfw.redhat.com started  
vm:guest6 botan.salab.dfw.redhat.com started  
root@ami:~  
File Edit View Terminal Help  
device-mapper: uevent: version 1.0.3  
device-mapper: ioctl: 4.11.5-ioclt (2007-12-12) initialised: dm-devel@redhat.com  
device-mapper: dm-raid45: initialized v0.2594l  
EXT3-fs: INFO: recovery required on readonly filesystem.  
EXT3-fs: write access will be enabled during recovery.  
kjournald starting. Commit interval 5 seconds  
EXT3-fs: recovery complete.  
EXT3-fs: mounted filesystem with ordered data mode.  
SELinux: Disabled at runtime.  
type=1404 audit(1251412033.542:2): selinux=0 auid=4294967295 ses=4294967295  
Bridge firewalling registered  
ip_tables: (C) 2000-2006 Netfilter Core Team  
Netfilter messages via NETLINK v0.30.  
ip_conntrack version 2.4 (8192 buckets, 65536 max) - 304 bytes per conntrack  
  
Red Hat Enterprise Linux Server release 5.4 (Tikanga)  
Kernel 2.6.18-164.el5xen on an x86_64  
  
sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump  
Kexec: Warning: crash image not loaded  
Kernel panic - not syncing: SysRq-triggered panic!  
  
[root@ami ~]#  
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]#
```

```
root@sayoko:~  
File Edit View Terminal Help  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 10:56:25 2009  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
Read from remote host sayoko.salab.dfw.redhat.com: Connection reset by peer  
Connection to sayoko.salab.dfw.redhat.com closed.  
[root@ai ~]#  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 17:23:27 2009 from ai.salab.dfw.redhat.com  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
[  
root@ayame:~  
File Edit View Terminal Help  
Every 2.0s: clustat Thu Aug 27 17:28:15 2009  
Cluster Status for dom0-cluster @ Thu Aug 27 17:28:15 2009  
Member Status: Quorate  
  
Member Name ID Status  
-----  
ami.salab.dfw.redhat.com 1 Online, rgmanager  
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com 3 Online, rgmanager  
  
Service Name Owner (Last) State  
-----  
vm:guest1 ayame.salab.dfw.redhat.com started  
vm:guest2 ami.salab.dfw.redhat.com started  
vm:guest3 ayame.salab.dfw.redhat.com started  
vm:guest4 ayame.salab.dfw.redhat.com started  
vm:guest5 botan.salab.dfw.redhat.com started  
vm:guest6 botan.salab.dfw.redhat.com started  
[  
root@ami:~  
File Edit View Terminal Help  
device-mapper: uevent: version 1.0.3  
device-mapper: ioctl: 4.11.5-ioclt (2007-12-12) initialised: dm-devel@redhat.com  
device-mapper: dm-raid45: initialized v0.2594l  
EXT3-fs: INFO: recovery required on readonly filesystem.  
EXT3-fs: write access will be enabled during recovery.  
kjournald starting. Commit interval 5 seconds  
EXT3-fs: recovery complete.  
EXT3-fs: mounted filesystem with ordered data mode.  
SELinux: Disabled at runtime.  
type=1404 audit(1251412033.542:2): selinux=0 auid=4294967295 ses=4294967295  
Bridge firewalling registered  
ip_tables: (C) 2000-2006 Netfilter Core Team  
Netfilter messages via NETLINK v0.30.  
ip_conntrack version 2.4 (8192 buckets, 65536 max) - 304 bytes per conntrack  
  
Red Hat Enterprise Linux Server release 5.4 (Tikanga)  
Kernel 2.6.18-164.el5xen on an x86_64  
  
sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump  
Kexec: Warning: crash image not loaded  
Kernel panic - not syncing: SysRq-triggered panic!  
  
[root@ami ~]# [  
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]# [  
[
```

Set up the clustered app on the domU cluster

Define a new cluster - domU-cluster

Cluster tab

Create new cluster

Cluster name

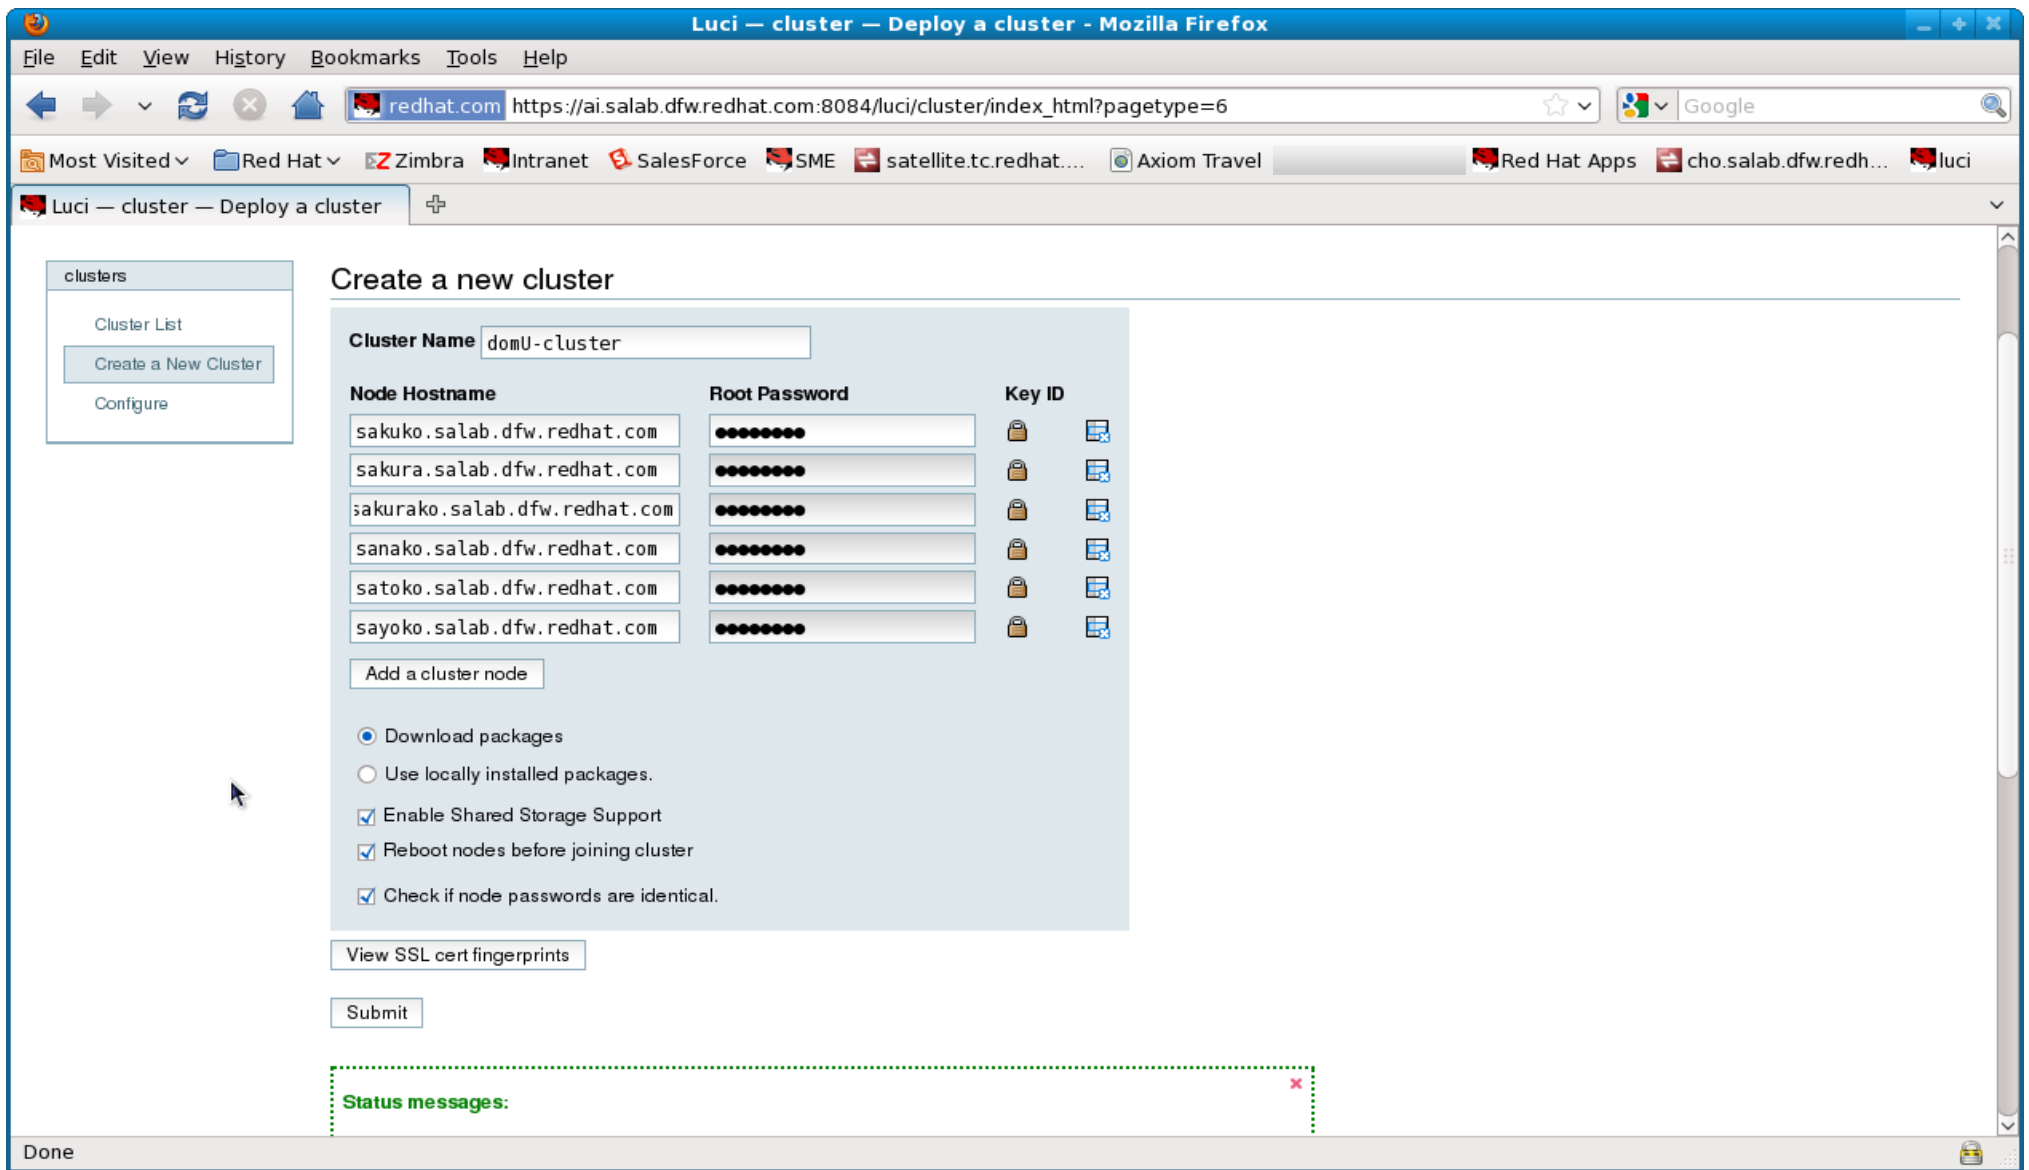
Node Hostnames/passwords

Enable share storage

Reboot nodes

Check if passwords identical

View SSL fingerprints



cluster — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

cluster — luci

redhat CLUSTER AND STORAGE SYSTEMS


homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure
- dom0-cluster
- domU-cluster


domU-cluster

Please be patient - this cluster's configuration is being modified.



Creating node "satoko.salab.dfw.redhat.com" for cluster "domU-cluster"
Node still being created


Install Reboot Configure Join



[Stop waiting for this job to complete](#)

Creating node "sakuko.salab.dfw.redhat.com" for cluster "domU-cluster"
Node still being created

Install Reboot Configure Join

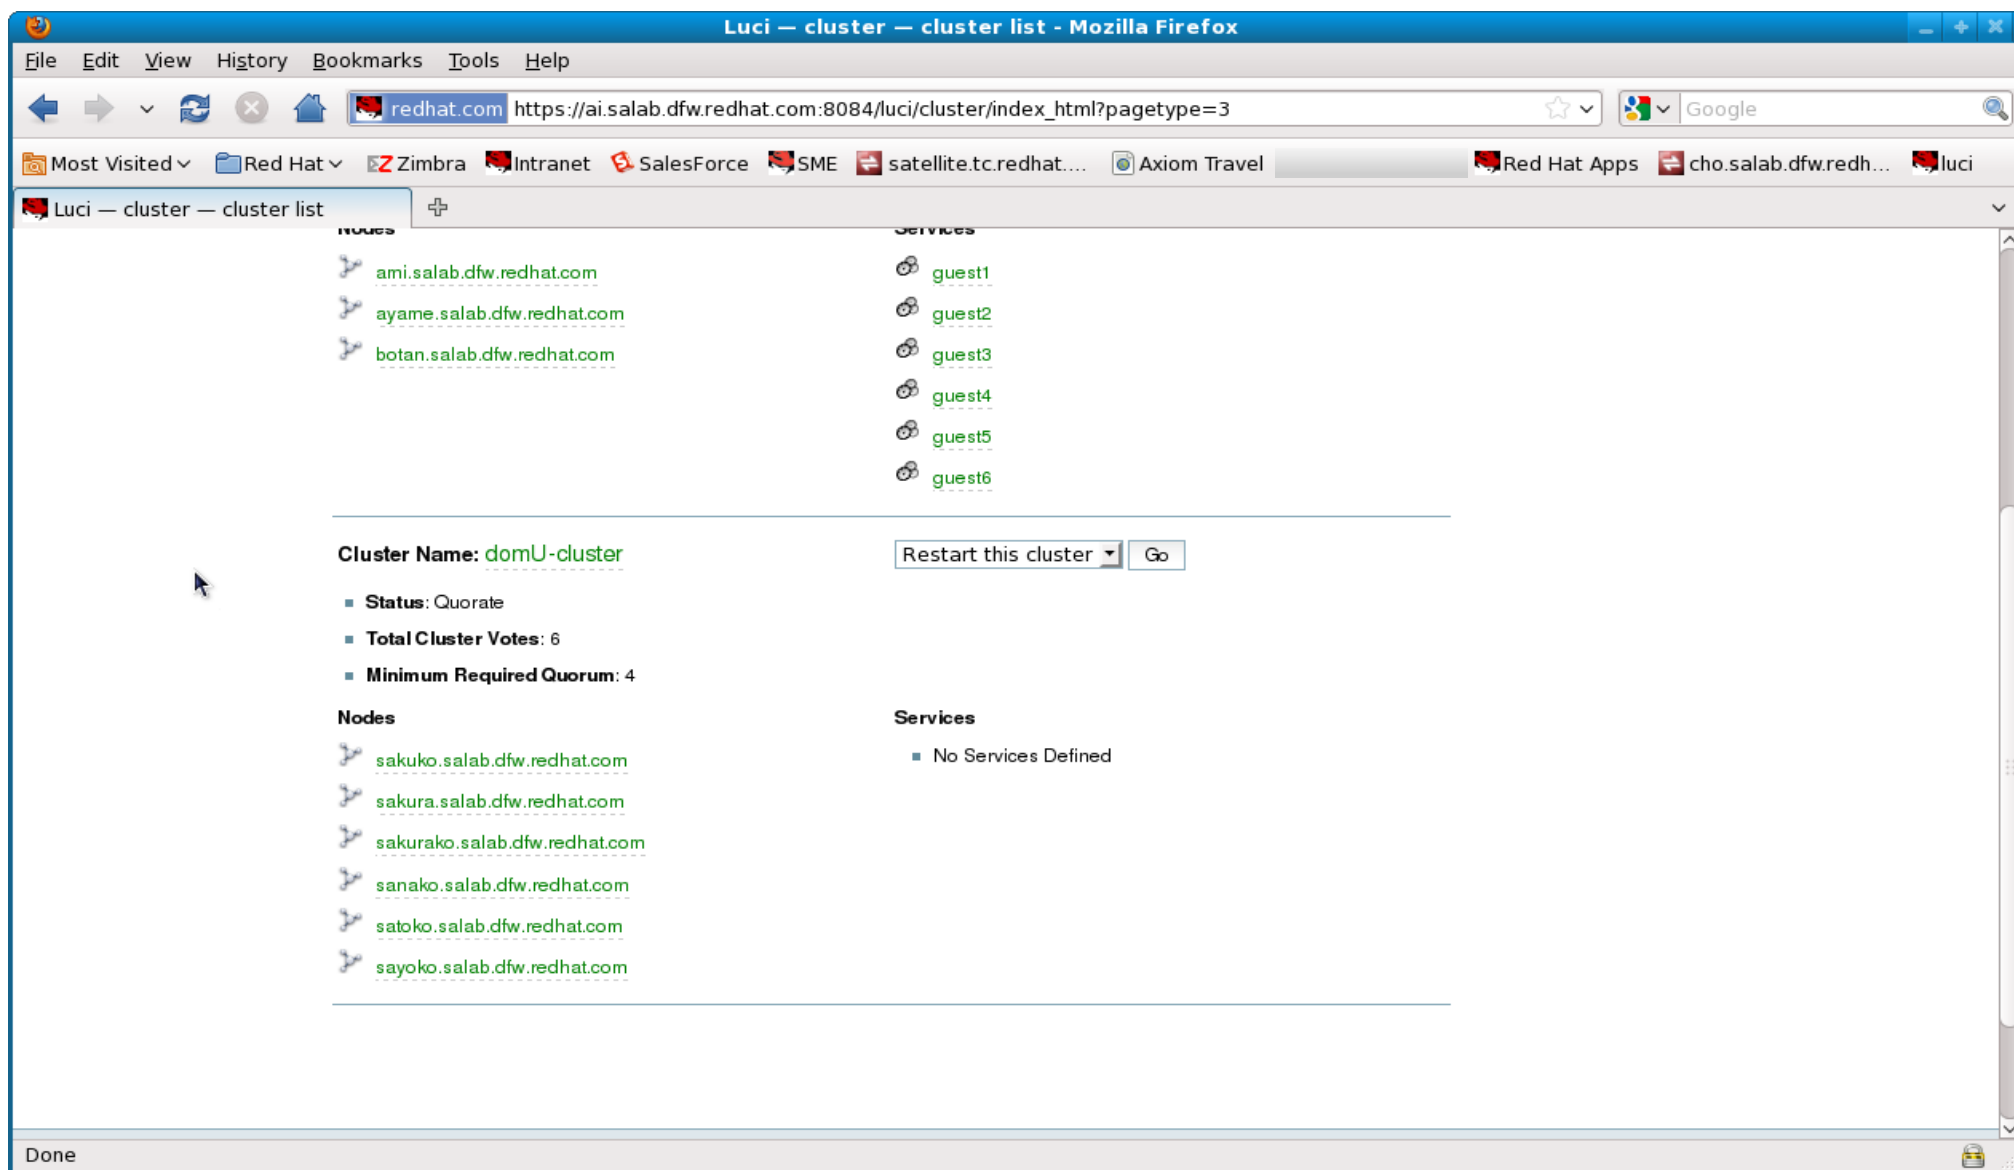


[Stop waiting for this job to complete](#)

Creating node "sakura.salab.dfw.redhat.com" for cluster "domU-cluster"
Node still being created

Transferring data from ai.salab.dfw.redhat.com...


```
root@ayame:~  
File Edit View Terminal Help  
vm:guest2          ami.salab.dfw.redhat.com      started  
vm:guest3          (none)                       recoverable  
vm:guest4          (none)                       recoverable  
vm:guest5          botan.salab.dfw.redhat.com    started  
vm:guest6          (none)                       recoverable  
[root@ayame ~]# clustat  
Cluster Status for dom0-cluster @ Fri Aug 28 13:56:35 2009  
Member Status: Quorate  
  
Member Name          ID    Status  
-----  
ami.salab.dfw.redhat.com      1 Online, rgmanager  
ayame.salab.dfw.redhat.com    2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com    3 Online, rgmanager  
  
Service Name          Owner (Last)          State  
-----  
vm:guest1          ami.salab.dfw.redhat.com      started  
vm:guest2          (none)                       recoverable  
vm:guest3          (none)                       recoverable  
vm:guest4          (none)                       recoverable  
vm:guest5          botan.salab.dfw.redhat.com    started  
vm:guest6          (none)                       recoverable  
[root@ayame ~]#
```



Set up the clustered app on the domU cluster

Copy the fence_xvm.key to the domU cluster members

Either scp or use the web UI

Set up the clustered app on the domU cluster

Set up a shared fence device

Cluster Tab

Cluster Name

Shared Fence Devices

Add a Fence Device

Type

Name

Submit and confirm

Luci — cluster — cluster list - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=3

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — cluster list Re: [Linux-cluster] how to us...

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

Choose a cluster to administer

Cluster Name: dom0-cluster Restart this cluster Go

- Status: Quorate
- Total Cluster Votes: 3
- Minimum Required Quorum: 2

Nodes

- ami.salab.dfw.redhat.com
- ayame.salab.dfw.redhat.com
- botan.salab.dfw.redhat.com

Services

- guest1
- guest2
- guest3
- guest4
- guest5
- guest6

Cluster Name: domU-cluster Restart this cluster Go

- Status: Quorate
- Total Cluster Votes: 6
- Minimum Required Quorum: 4

Nodes

- sakuko.salab.dfw.redhat.com

Services

- No Services Defined

https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=domU-cluster

Luci — cluster — Configure cluster properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=7&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — Configure cl... Re: [Linux-cluster] how to us...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure
 - dom0-cluster
 - domU-cluster

domu-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence Devices

domU-cluster

General Fence Multicast Quorum Partition

General Properties

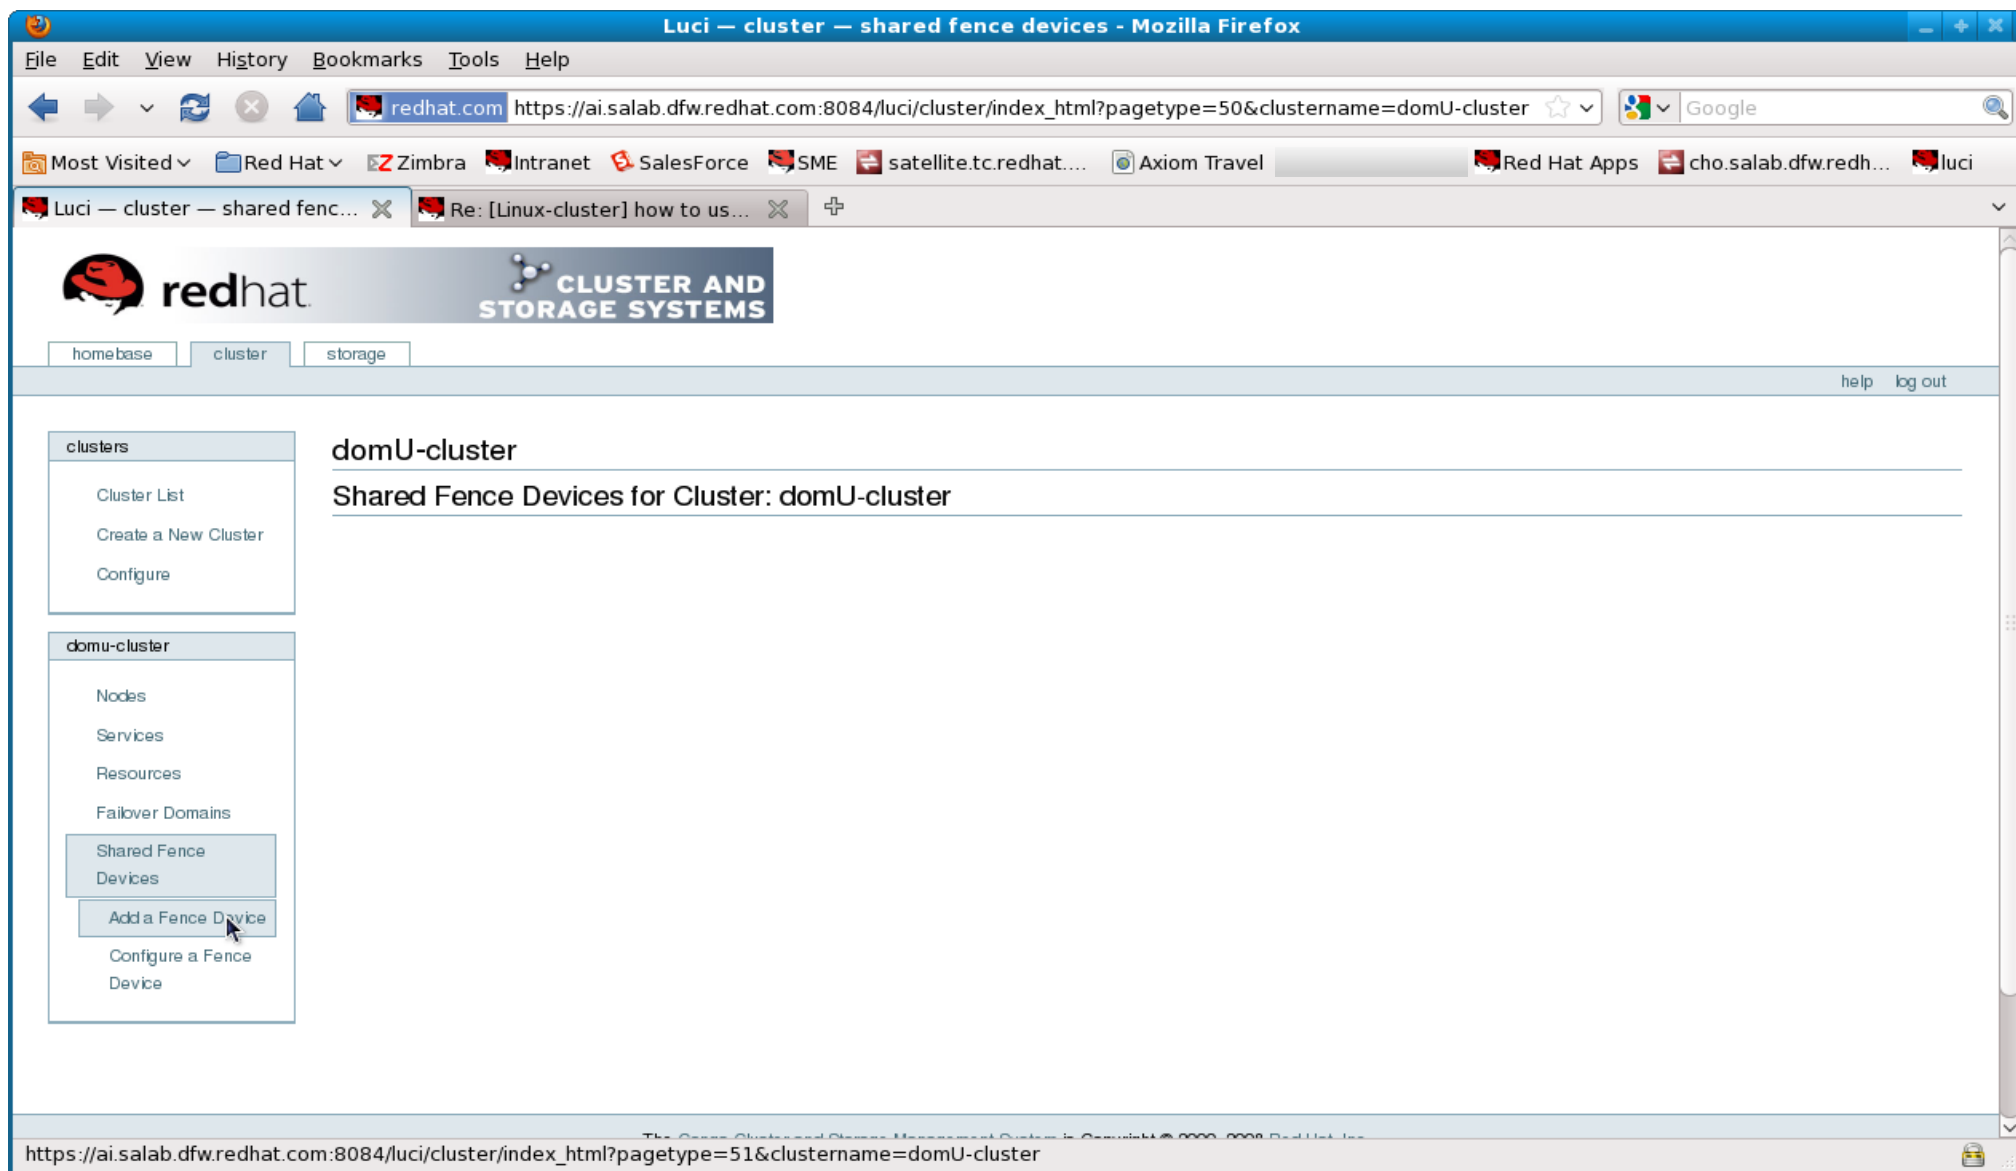
Cluster Name

Configuration Version

► Show advanced cluster properties

The Conga Cluster and Storage Management System is Copyright © 2000–2008 Red Hat, Inc.

https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=50&clustername=domU-cluster



Luci — cluster — fence devices - Add a new fence device - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=51&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — fence devices - ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
- Failover Domains
- Shared Fence Devices
- Add a Fence Device
- Configure a Fence Device

domU-cluster

Add a Sharable Fence Device

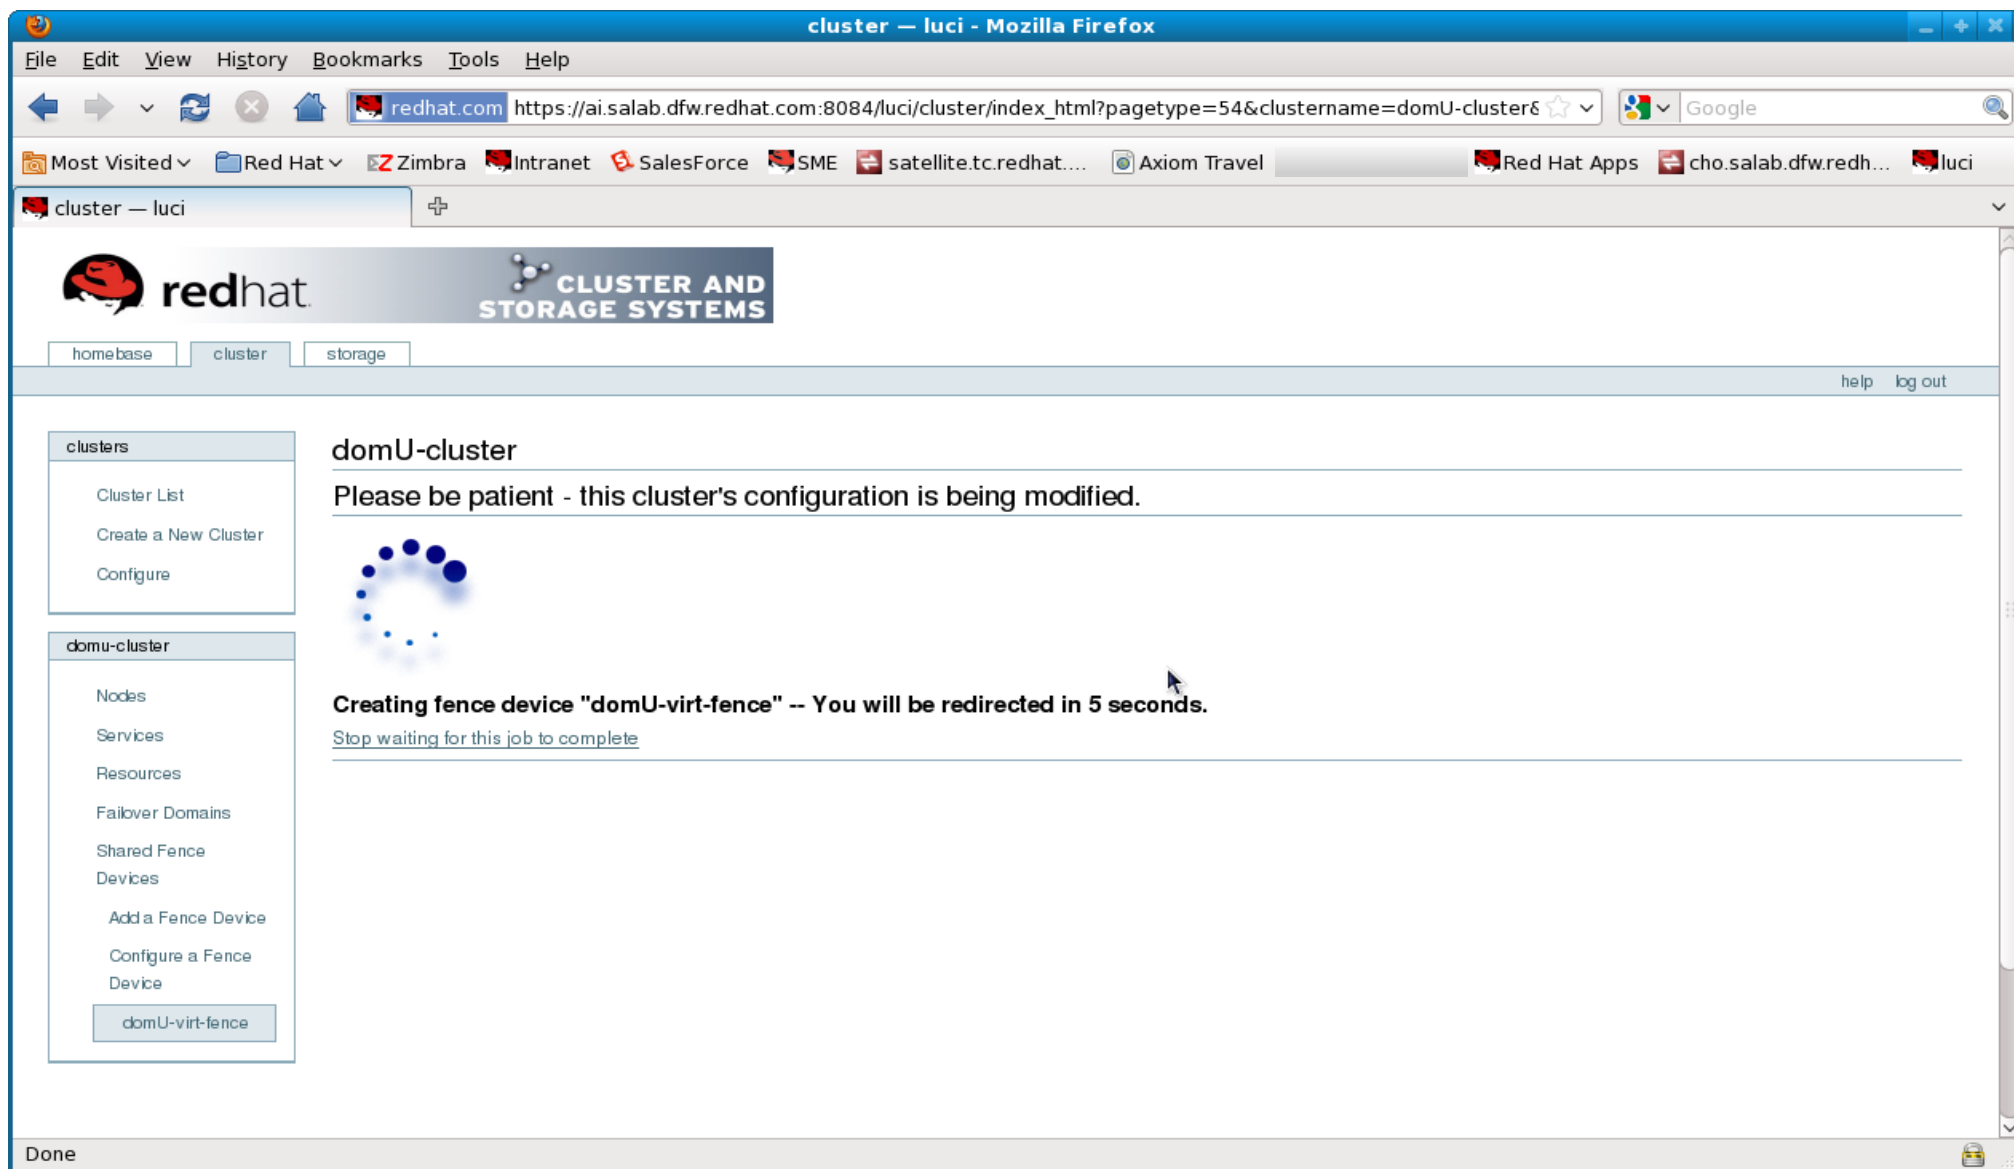
Fencing Type

Virtual Machine Fencing

Fence Type Virtual Machine Fencing

Name

Done



Set up the clustered app on the domU cluster

Assign Fence Devices to Each domU

Cluster Node

Main Fencing Method

Add Fence Device To This Level

Pre-defined device

Domain name

Update and confirm

Repeat for each domU machine

Luci — cluster — nodes — properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&clustername=domU-cluster&r

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — nodes — proper... +

Cluster List
Create a New Cluster
Configure

domu-cluster

Nodes
Add a Node
Configure
satoko.salab.dfw.redhat.com
sakuko.salab.dfw.redhat.com
sakura.salab.dfw.redhat.com
sayoko.salab.dfw.redhat.com
sanako.salab.dfw.redhat.com
sakurako.salab.dfw.redhat.com
Services
Resources
Failover Domains
Shared Fence Devices

Node Name: sakuko.salab.dfw.redhat.com

Choose a Task... Go

Status: Cluster member

Show recent log activity for this node

Cluster daemons running on this node

| Daemon | Currently running | Enabled at start-up |
|-----------|-------------------|-------------------------------------|
| cman | yes | <input checked="" type="checkbox"/> |
| rgmanager | yes | <input checked="" type="checkbox"/> |

Update node daemon properties

Services on this Node

- No cluster services are currently running here

Failover Domain Membership

- This node has no failover domain membership

Main Fencing Method

Add a fence device to this level

Update main fence properties

Backup Fencing Method

Add a fence device to this level

Update backup fence properties

Done

Luci — cluster — nodes — properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&nodename=sayoko.salab.dfw. Google

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — nodes — proper...

Nodes

- Add a Node
- Configure
 - satoko.salab.dfw.redhat.com
 - sakuko.salab.dfw.redhat.com
 - sakura.salab.dfw.redhat.com
 - sayoko.salab.dfw.redhat.com**
 - sanako.salab.dfw.redhat.com
 - sakurako.salab.dfw.redhat.com
- Services
- Resources
- Failover Domains
- Shared Fence Devices

Cluster daemons running on this node

| Daemon | Currently running | Enabled at start-up |
|-----------|-------------------|-------------------------------------|
| cman | yes | <input checked="" type="checkbox"/> |
| rgmanager | yes | <input checked="" type="checkbox"/> |

[Update node daemon properties](#)

Services on this Node

- No cluster services are currently running here

Failover Domain Membership

- This node has no failover domain membership

Main Fencing Method

Fence Type Virtual Machine Fencing

Name

Domain

[Remove this instance](#)

[Remove this device](#) [Add an instance](#)

[Add a fence device to this level](#)

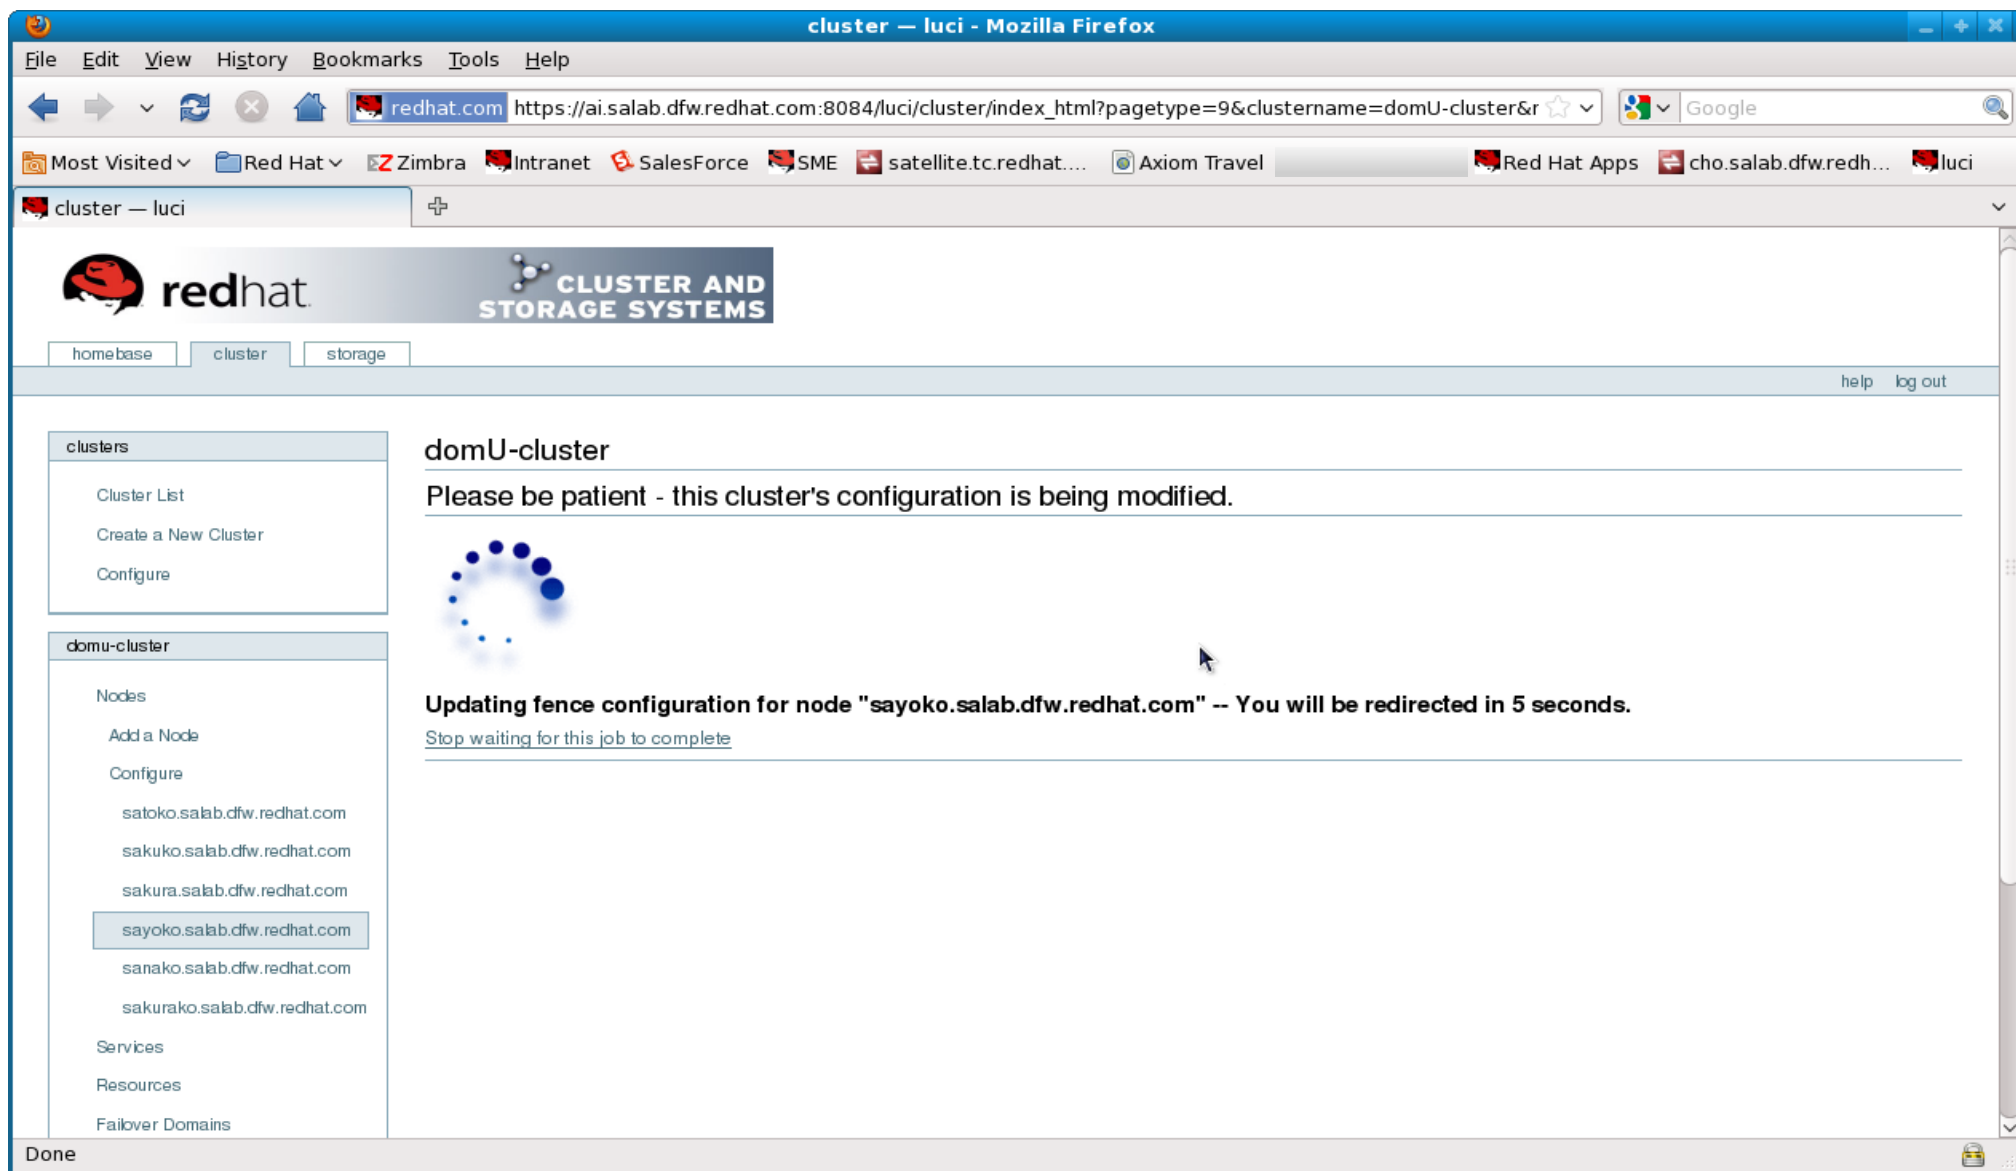
[Update main fence properties](#)

Backup Fencing Method

[Add a fence device to this level](#)

[Update backup fence properties](#)

Done



Set up the clustered app on the domU cluster

Test by fencing a domU

Luci — cluster — nodes — properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&nodename=satoko.salab.dfw

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — nodes — proper...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

Nodes

- Add a Node
- Configure
- satoko.salab.dfw.redhat.com
- sakuko.salab.dfw.redhat.com
- sakura.salab.dfw.redhat.com
- sayoko.salab.dfw.redhat.com
- sanako.salab.dfw.redhat.com
- sakurako.salab.dfw.redhat.com

Services

Resources

Failover Domains

domU-cluster

Node Name: satoko.salab.dfw.redhat.com

Status: Cluster member

Show recent log activity for this node

Fence this node Go

Cluster daemons running on this node

| Daemon | Currently running | Enabled at start-up |
|-----------|-------------------|-------------------------------------|
| cman | yes | <input checked="" type="checkbox"/> |
| rgmanager | yes | <input checked="" type="checkbox"/> |

Update node daemon properties

Services on this Node

- No cluster services are currently running here

Failover Domain Membership

- This node has no failover domain membership

Done

cluster — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=10&clustername=domU-cluster6

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

cluster — luci Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure


domu-cluster

Nodes

- Add a Node
- Configure
 - satoko.salab.dfw.redhat.com
 - sakuko.salab.dfw.redhat.com
 - sakura.salab.dfw.redhat.com
 - sayoko.salab.dfw.redhat.com
 - sanako.salab.dfw.redhat.com
 - sakurako.salab.dfw.redhat.com
- Services
- Resources
- Failover Domains

domU-cluster

Please be patient - this cluster's configuration is being modified.



Node "sakuko.salab.dfw.redhat.com" is being fenced by node "satoko.salab.dfw.redhat.com" -- You will be redirected in 5 seconds.

[Stop waiting for this job to complete](#)

Done


```
root@satoko:~  
File Edit View Terminal Help  
Aug 28 15:00:11 satoko fence_node[3646]: Fence of "sakuko.salab.dfw.redhat.com"  
was successful  
Aug 28 15:00:20 satoko openais[2203]: [TOTEM] entering GATHER state from 12.  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] entering GATHER state from 0.  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] Saving state aru 91 high seq recei  
ved 91  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] Storing new sequence id for ring 6  
c  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] entering COMMIT state.  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] entering RECOVERY state.  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] position [0] member 10.15.183.92:  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] previous ring seq 104 rep 10.15.18  
3.91  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] aru 91 high delivered 91 received  
flag 1  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] position [1] member 10.15.183.93:  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] previous ring seq 104 rep 10.15.18  
3.91  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] aru 91 high delivered 91 received  
flag 1  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] position [2] member 10.15.183.94:  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] previous ring seq 104 rep 10.15.18  
3.91
```

Set up the clustered app on the domU cluster

Define failover domain(s) for the app you want to run on this domU cluster

In this case, a simple Apache web server with an NFS mounted `/var/www/html` and a floating IP address

Here I've defined a FD `prefer_satoko`

Luci — cluster — failover domains — Add a failover domain - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=41&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — failover do... Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
- Failover Domains
 - Add a Failover Domain
 - Configure a Failover Domain
- Shared Fence Devices

domU-cluster

Add a Failover Domain

Failover Domain Name

Prioritized ☒

Restrict failover to this domain's members ☒

Do not fail back services in this domain ☐

Failover domain membership

| Node | Member | Priority |
|-------------------------------|-------------------------------------|---------------------------------|
| satoko.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | <input type="text" value="1"/> |
| sakuko.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | <input type="text" value="10"/> |
| sakura.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | <input type="text" value="20"/> |
| sayoko.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | <input type="text" value="30"/> |
| sanako.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | <input type="text" value="40"/> |
| sakurako.salab.dfw.redhat.com | <input checked="" type="checkbox"/> | <input type="text" value="50"/> |

Done

Set up the clustered app on the domU cluster

Define resources for the service

- Cluster Tab

- Cluster Name

- Resources

- Add a Resource

 - IP address

 - NFS mount

 - Script

Luci — cluster — resources — Add a resource - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=31&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — resources ... Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
- Add a Resource**
- Configure a Resource
- Failover Domains
- Shared Fence Devices

domU-cluster

Add a Resource

Select a Resource Type

Select a resource type

Add a Resource to this cluster

https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=31&clustername=domU-cluster

Luci — cluster — resources — Add a resource - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=31&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — resources ... Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
 - Add a Resource
 - Configure a Resource
- Failover Domains
- Shared Fence Devices

domU-cluster

Add a Resource

IP Address Resource Configuration

IP address

Monitor link ☒

Done

Luci — cluster — resources - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=30&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — resources Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
- Add a Resource
- Configure a Resource
- Failover Domains
- Shared Fence Devices

domU-cluster

Resources for domu-cluster

| Resource Name | Type | Configure | Delete |
|---------------|------------|---------------------------|------------------------|
| 10.15.183.248 | IP Address | configure | delete |

Done

The Open Clustered Storage Management System is Copyright © 2000-2008 Red Hat, Inc.

Luci — cluster — resources — Add a resource - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=31&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — resources ... Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
 - Add a Resource
 - Configure a Resource
- Failover Domains
- Shared Fence Devices

domU-cluster

Add a Resource

NFS Mount Resource Configuration

Name: web-mount

Mount point: /var/www/html

Host: salab.dfw.redhat.com

Export path: /var/nfs

NFS version: ☒ NFS3 ☐ NFS4

Options: ro,soft,intr,rsize=3

Force unmount: ☐

Submit

Done

Luci — cluster — resources - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=30&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — resources Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
- Add a Resource
- Configure a Resource
- Failover Domains
- Shared Fence Devices

domU-cluster

Resources for domu-cluster

| Resource Name | Type | Configure | Delete |
|---------------|------------|---------------------------|------------------------|
| 10.15.183.248 | IP Address | configure | delete |
| web-mount | NFS Mount | configure | delete |

Done

Luci — cluster — resources — Add a resource - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=31&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — resources ... Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
- Add a Resource**
- Configure a Resource
- Failover Domains
- Shared Fence Devices

domU-cluster

Add a Resource

Script Resource Configuration

Name

Full path to script file

Done

Luci — cluster — resources - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=30&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — resources Bug 435189 - fenced admin ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Resources
- Add a Resource
- Configure a Resource
- Failover Domains
- Shared Fence Devices

domU-cluster

Resources for domu-cluster

| Resource Name | Type | Configure | Delete |
|---------------|------------|---------------------------|------------------------|
| 10.15.183.248 | IP Address | configure | delete |
| web-mount | NFS Mount | configure | delete |
| httpd-script | Script | configure | delete |

Done

Set up the clustered app on the domU cluster

Define a service which uses all of those resources

Cluster Tab

Cluster Name

Services

Set up the clustered app on the domU cluster

Define a service which uses all of those resources

Add a Service

Name

Auto-start?

Enable NFS lock workarounds

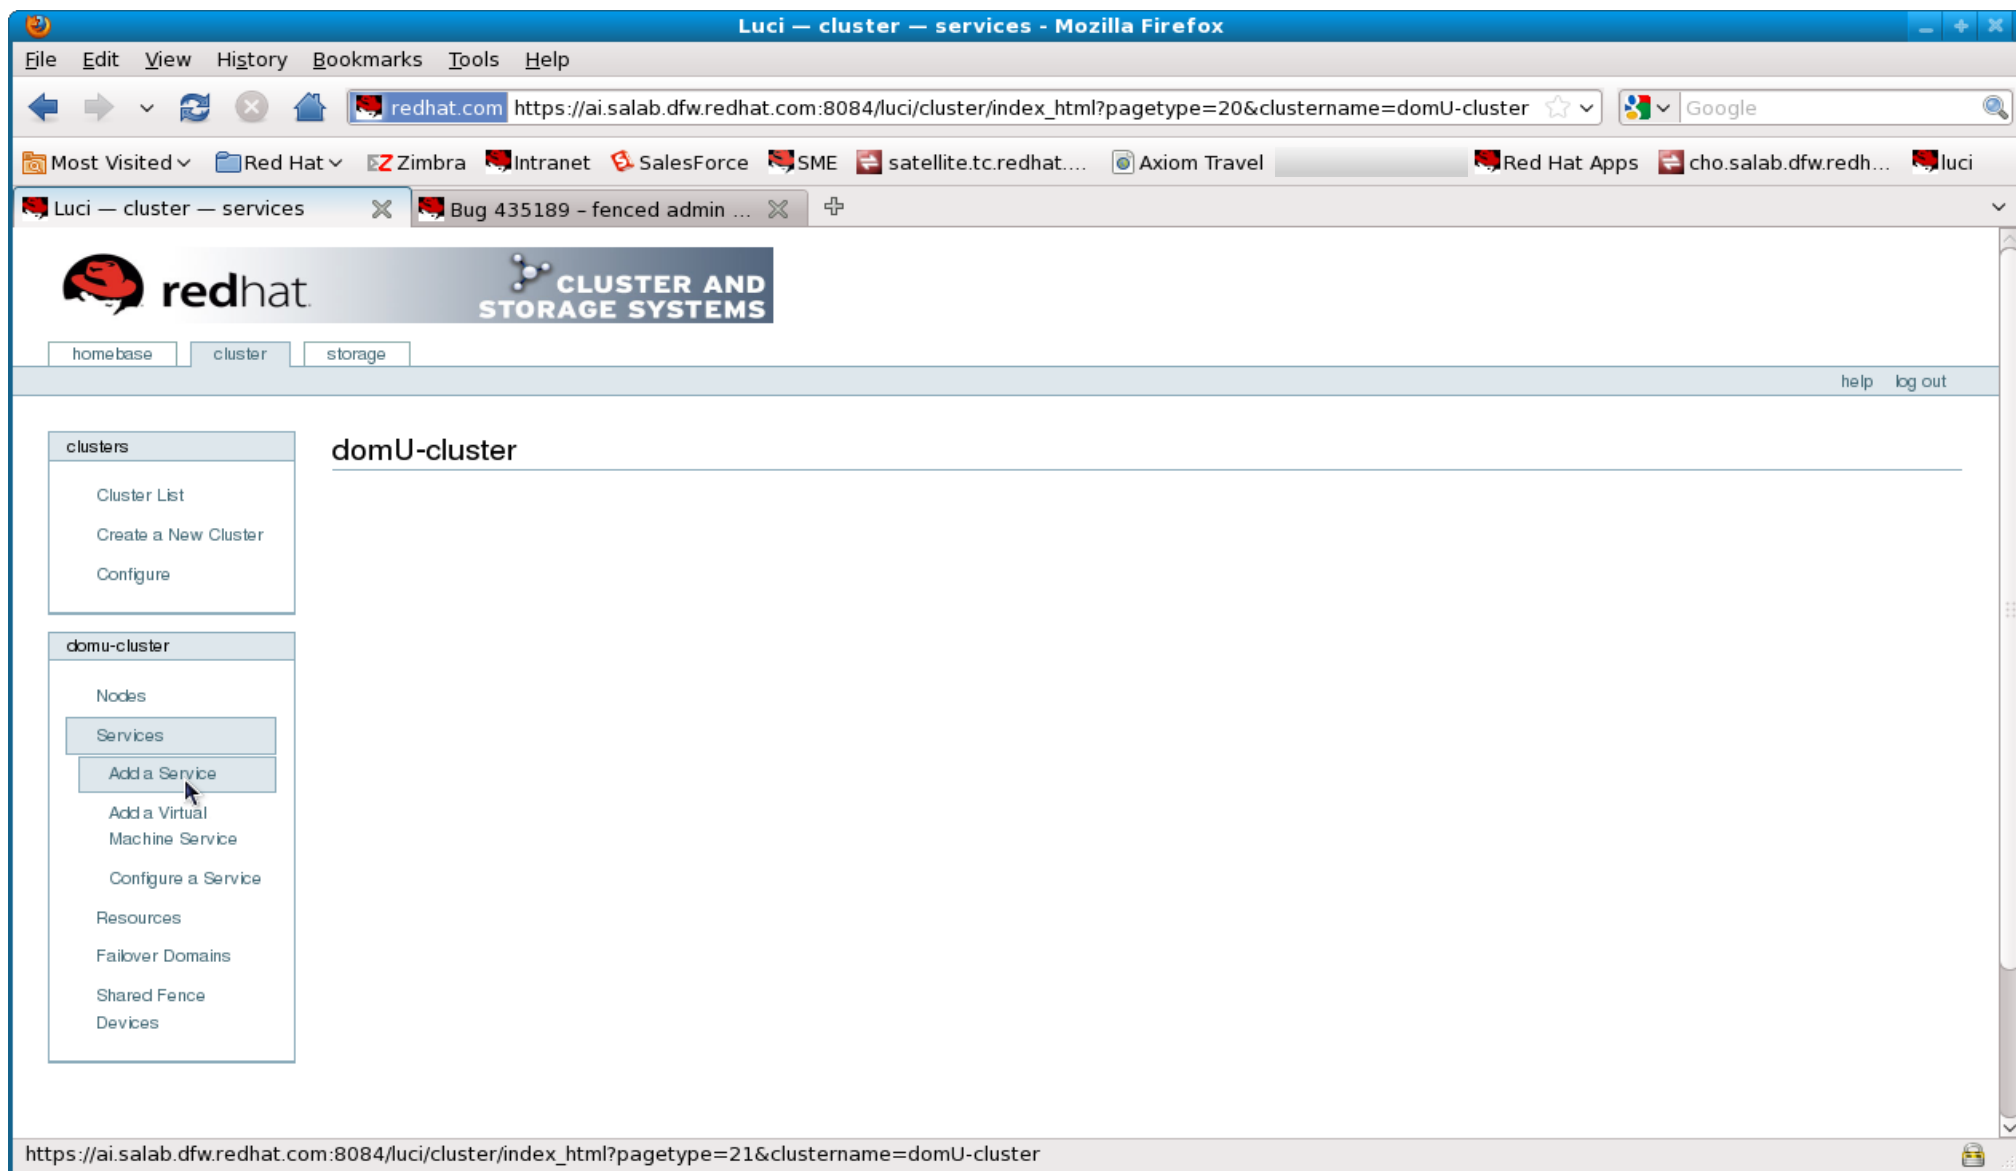
Run exclusive

FD

Recovery Policy

Maximum number of restart failures before relocating

Length of time in seconds after which to forget a restart



Luci — cluster — services — Add a new service - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=21&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services — Add ...

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
 - Add a Service
 - Add a Virtual Machine Service
 - Configure a Service
- Resources
- Failover Domains
- Shared Fence
- Devices

domU-cluster

Add a Service

Service name web-service

Automatically start this service ☒

Enable NFS lock workarounds ☐

Run exclusive ☐

Failover Domain prefer_satoko

Recovery policy Restart

Maximum number of restart failures before relocating 1

Length of time in seconds after which to forget a restart 600

Add a resource to this service Submit

Find: workaro Previous Next Highlight all Match case

Done

Set up the clustered app on the domU cluster

Define a service which uses all of those resources

Add Global Resources to the service

IP address

NFS Mount

Script

Luci — cluster — services — Add a new service - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=21&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services — Add ...

Cluster List
Create a New Cluster
Configure

domu-cluster

Nodes
Services
Add a Service
Add a Virtual Machine Service
Configure a Service
Resources
Failover Domains
Shared Fence
Devices

Add a Service

Service name web-service

Automatically start this service ☒

Enable NFS lock workarounds ☒

Run exclusive ☐

Failover Domain prefer_satoko

Recovery policy Restart

Maximum number of restart failures before relocating 1

Length of time in seconds after which to forget a restart 600

Add a new local resource

Select a resource type

or

Use an existing global resource

Select a resource name

Add a resource to this service Submit

Find: workaro Previous Next Highlight all Match case

Done

Luci — cluster — services — Add a new service - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=21&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services — Add ...

Cluster List
Create a New Cluster
Configure

domu-cluster

Nodes
Services
Add a Service
Add a Virtual Machine Service
Configure a Service
Resources
Failover Domains
Shared Fence
Devices

Add a Service

Service name web-service

Automatically start this service ☒

Enable NFS lock workarounds ☒

Run exclusive ☐

Failover Domain prefer_satoko

Recovery policy Restart

Maximum number of restart failures before relocating 1

Length of time in seconds after which to forget a restart 600

IP Address Resource Configuration

IP address 10.15.183.248

Monitor link ☒

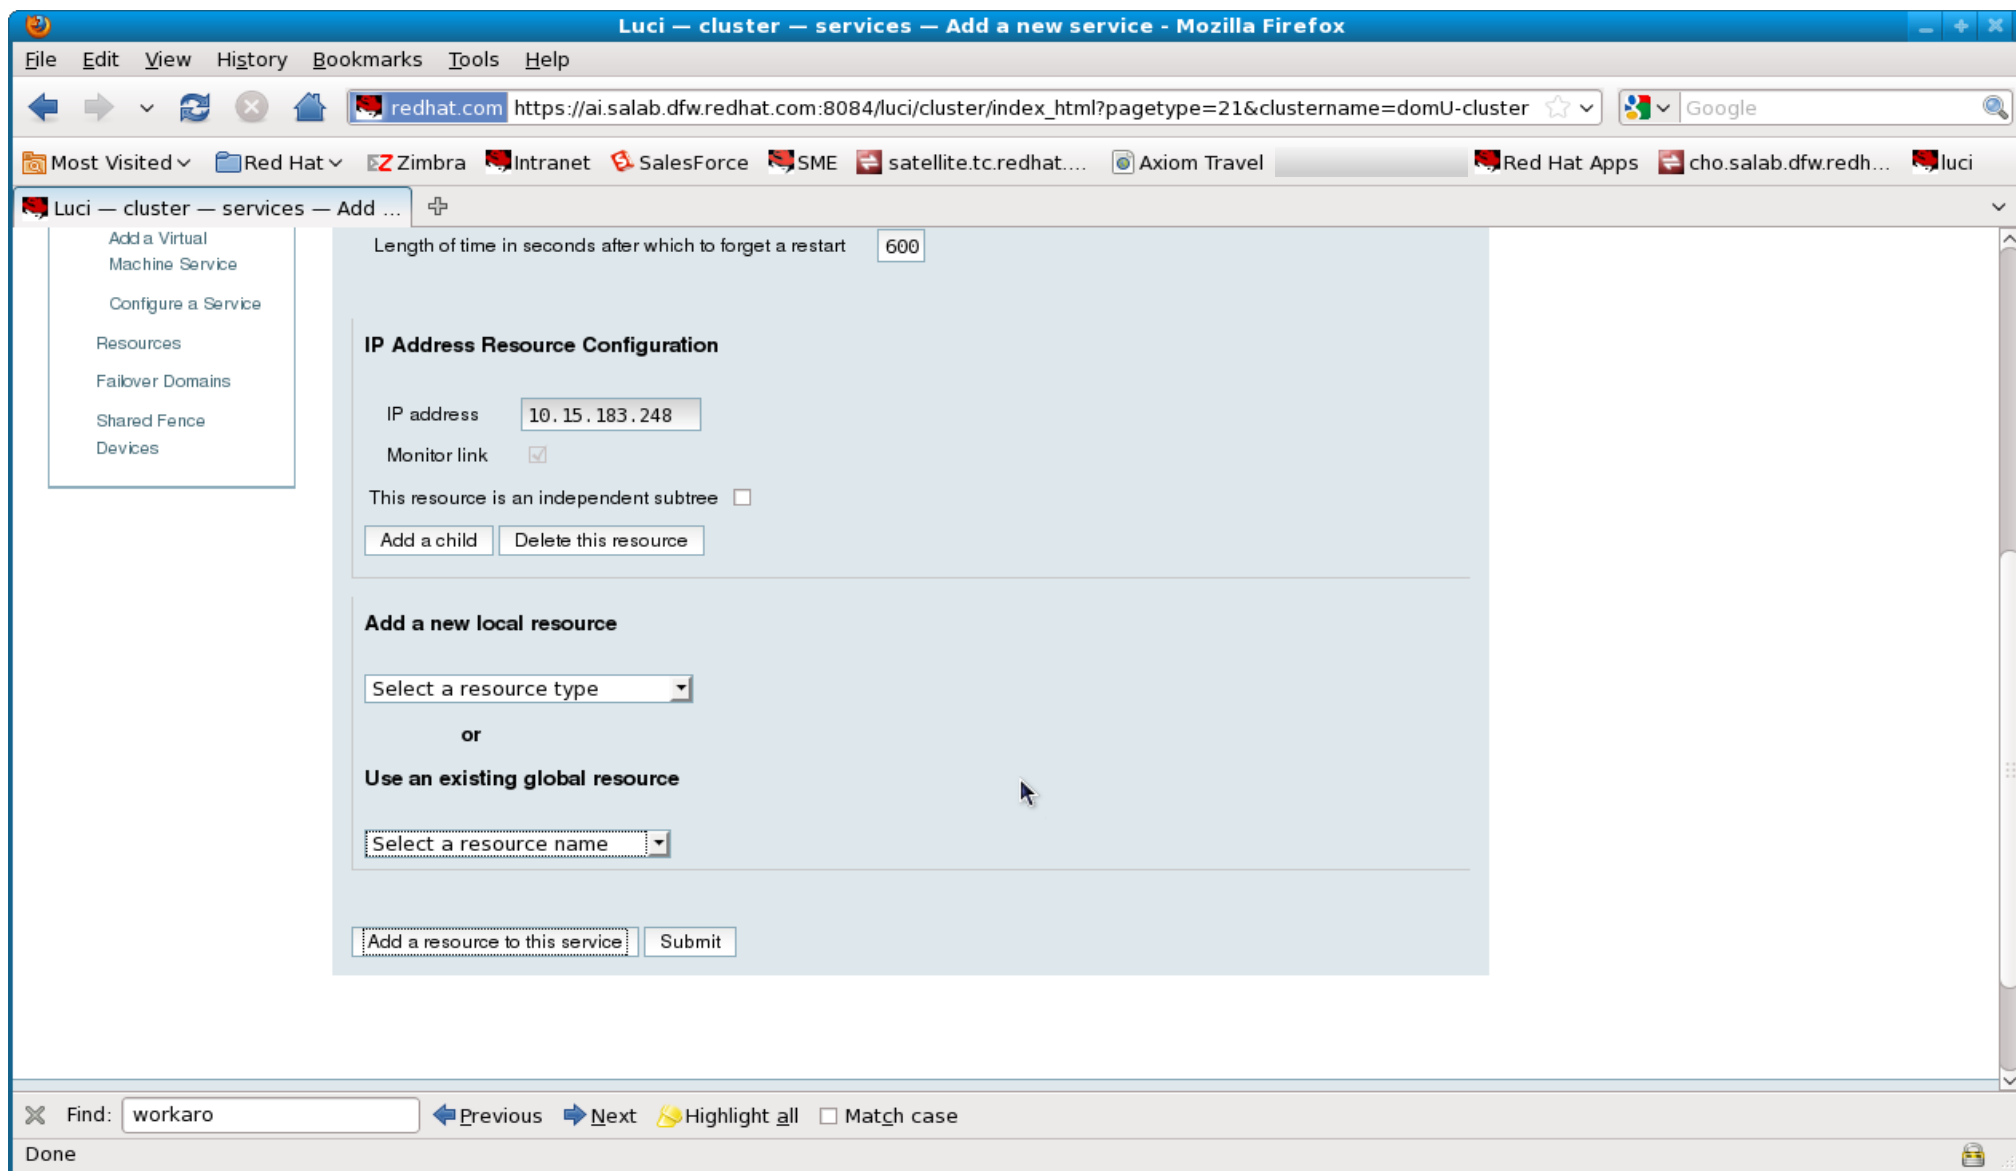
This resource is an independent subtree ☐

Add a child Delete this resource

Add a resource to this service Submit

Find: workaro Previous Next Highlight all Match case

Done



Luci — cluster — services — Add a new service - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=21&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services — Add ...

- Resources
- Failover Domains
- Shared Fence
- Devices

IP Address Resource Configuration

IP address: 10.15.183.248

Monitor link: ☒

This resource is an independent subtree: ☐

Add a child Delete this resource

NFS Mount Resource Configuration

Name: web-mount

Mount point: /var/www/html

Host: ai.salab.dfw.redhat..

Export path: /var/nfs

NFS version: ☒ NFS3 ☐ NFS4

Options: ro,soft,intr,rsize=3

Force unmount: ☐

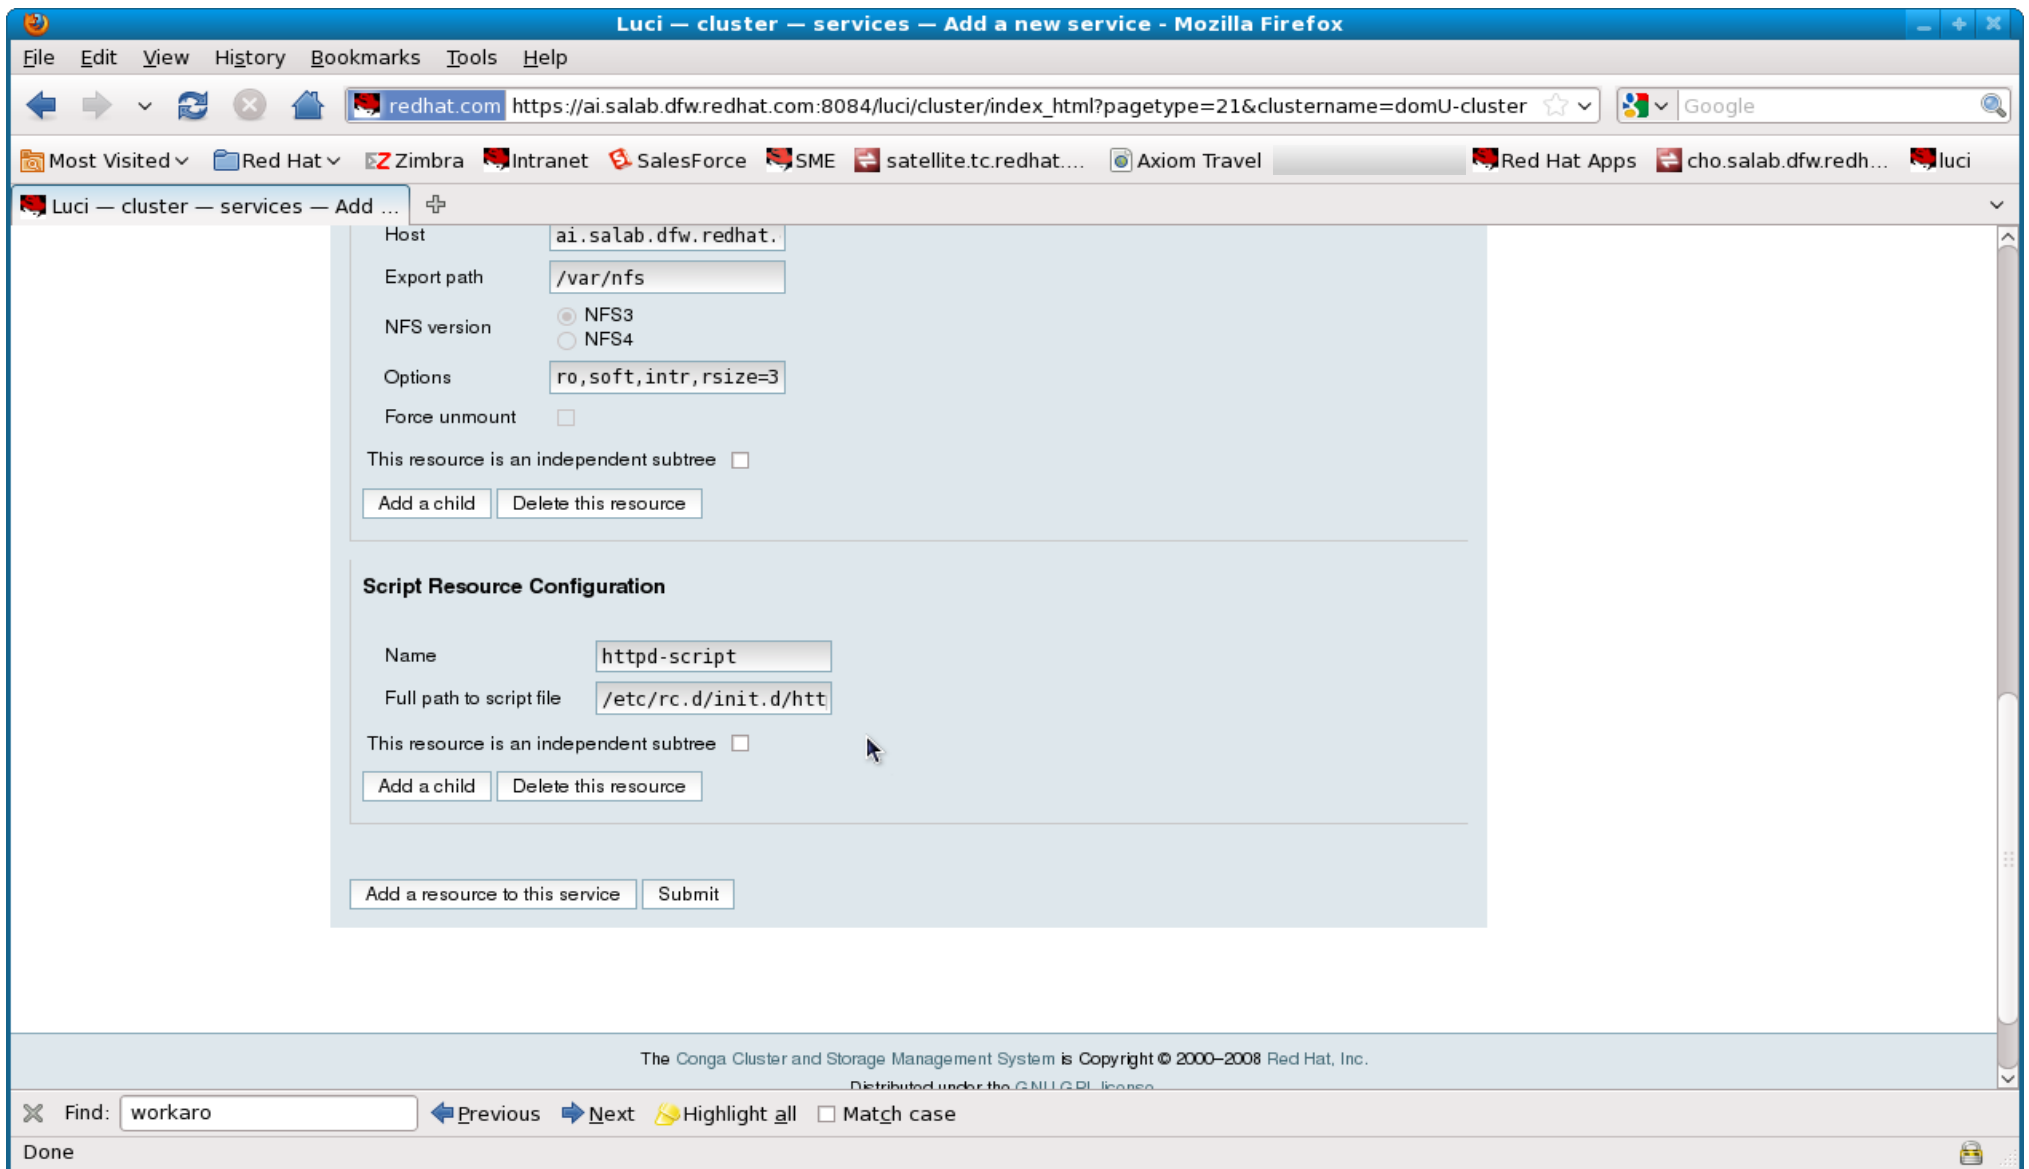
This resource is an independent subtree: ☐

Add a child Delete this resource

Add a resource to this service Submit

Find: workaro Previous Next Highlight all Match case

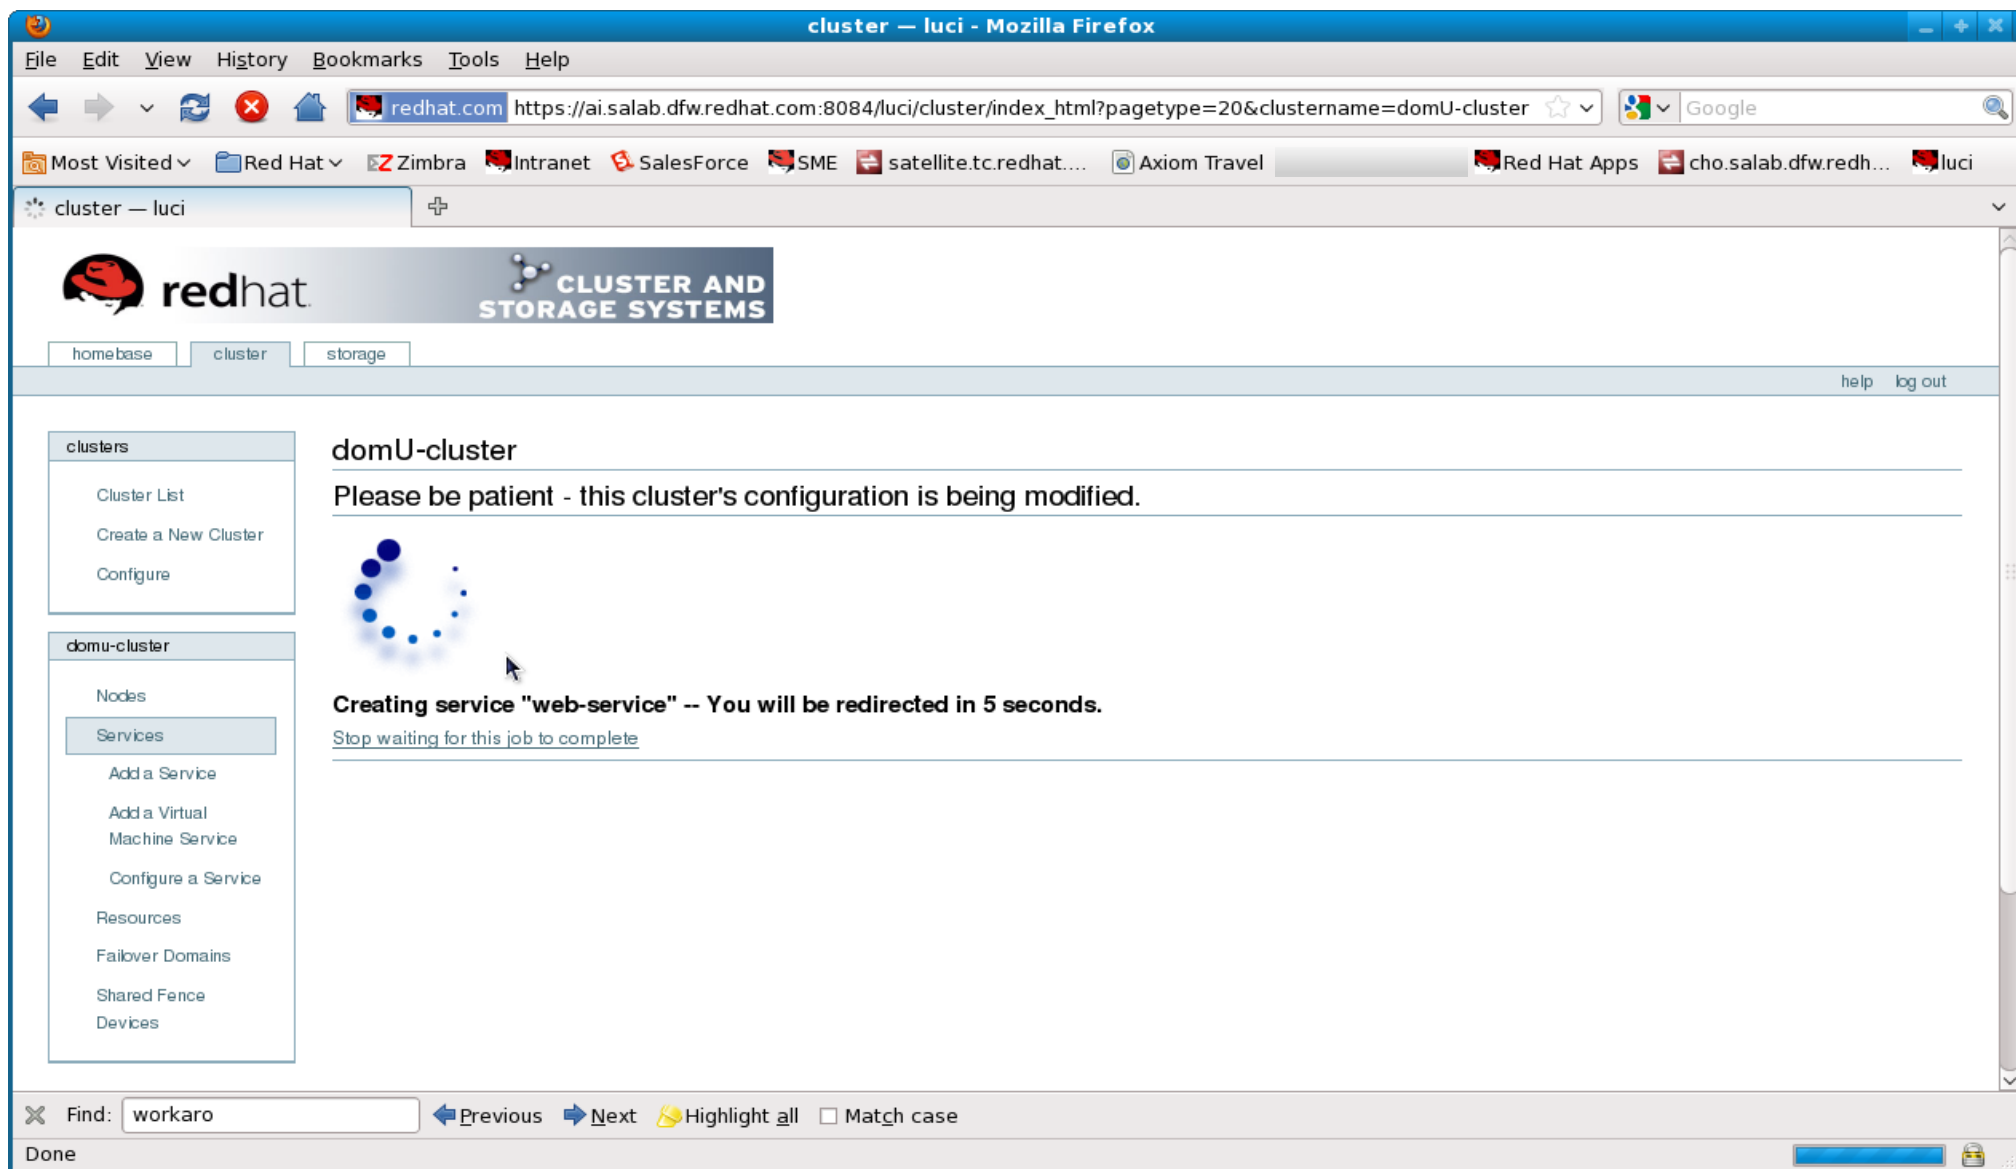
Done



Set up the clustered app on the domU cluster

Submit and confirm

Open the IP address defined in the resource definition in your web browser



Luci — cluster — services - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=20&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Add a Service
- Add a Virtual Machine Service
- Configure a Service
- Resources
- Failover Domains
- Shared Fence
- Devices

domU-cluster

Service Name **web-service**


Status: Running on satoko.salab.dfw.redhat.com
Autostart is enabled for this service

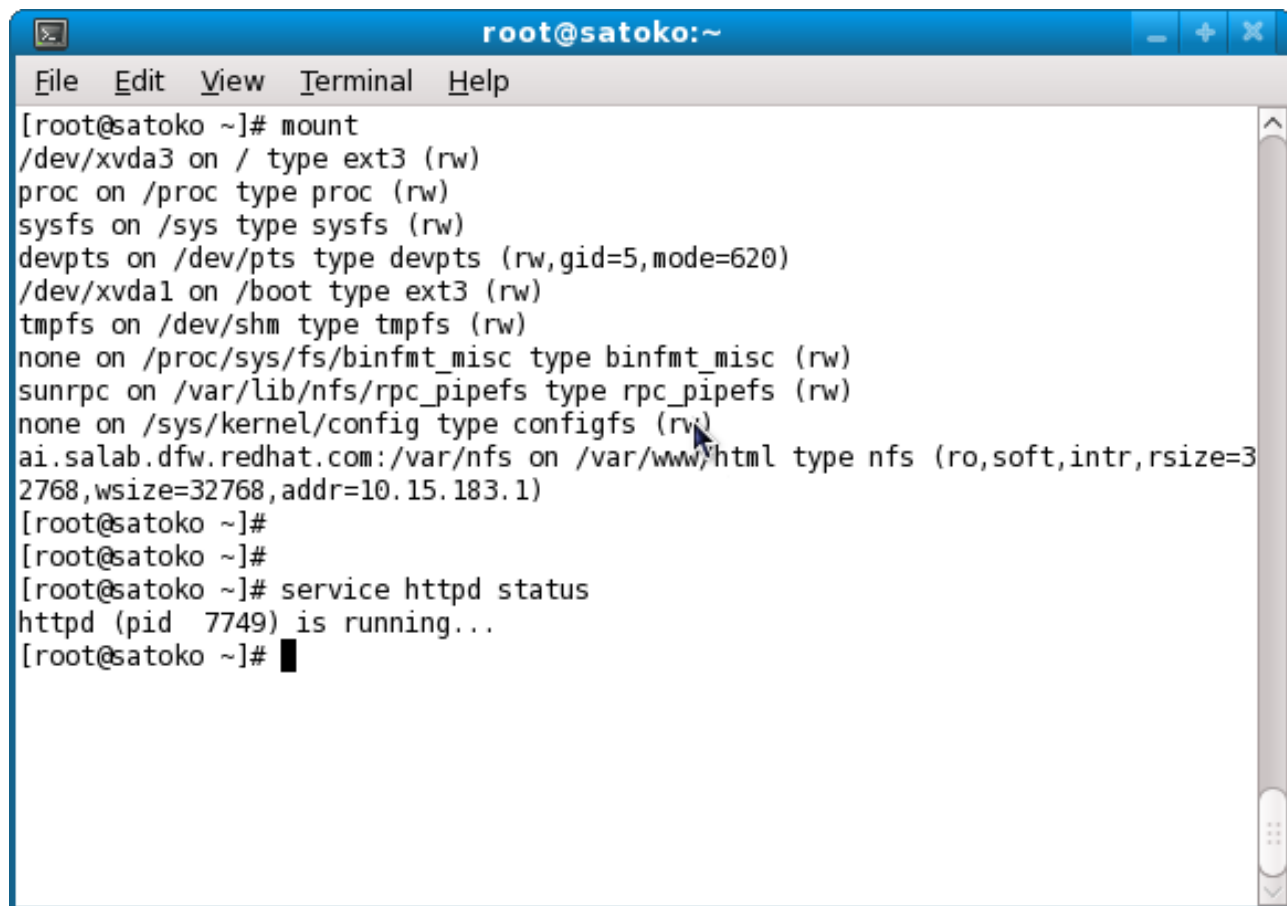
Failover Domain Association: prefer_satoko

Choose a Task... Go

Find: workaro Previous Next Highlight all Match case

Done





```
root@satoko:~  
File Edit View Terminal Help  
[root@satoko ~]# mount  
/dev/xvda3 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/xvda1 on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /sys/kernel/config type configfs (rw)  
ai.salab.dfw.redhat.com:/var/nfs on /var/www/html type nfs (ro,soft,intr,rsize=32768,wsiz  
e=32768,addr=10.15.183.1)  
[root@satoko ~]#  
[root@satoko ~]#  
[root@satoko ~]# service httpd status  
httpd (pid 7749) is running...  
[root@satoko ~]#
```

Set up the clustered app on the domU cluster

Test migration and failover

Luci — cluster — services - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=20&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Add a Service
- Add a Virtual Machine Service
- Configure a Service
- Resources
- Failover Domains
- Shared Fence
- Devices

domU-cluster

Service Name web-service

Status: Running on satoko.salab.dfw.redhat.com
Autostart is enabled for this service

Failover Domain Association: prefer_satoko

Relocate this service to sakuko.salab.dfw.redhat.com Go

Find: workaro Previous Next Highlight all Match case

Done

cluster — luci - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=22&clustername=domU-cluster6

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

cluster — luci

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters


- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Add a Service
- Add a Virtual Machine Service
- Configure a Service
- Resources
- Failover Domains
- Shared Fence
- Devices

domU-cluster

Please be patient - this cluster's configuration is being modified.



Starting cluster service "web-service" on node "sakuko.salab.dfw.redhat.com"

[Stop waiting for this job to complete](#)

Find: workaro Previous Next Highlight all Match case

Done

Luci — cluster — services - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=22&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domu-cluster

- Nodes
- Services
- Add a Service
- Add a Virtual Machine Service
- Configure a Service
- Resources
- Failover Domains
- Shared Fence
- Devices

domU-cluster

Service Name web-service

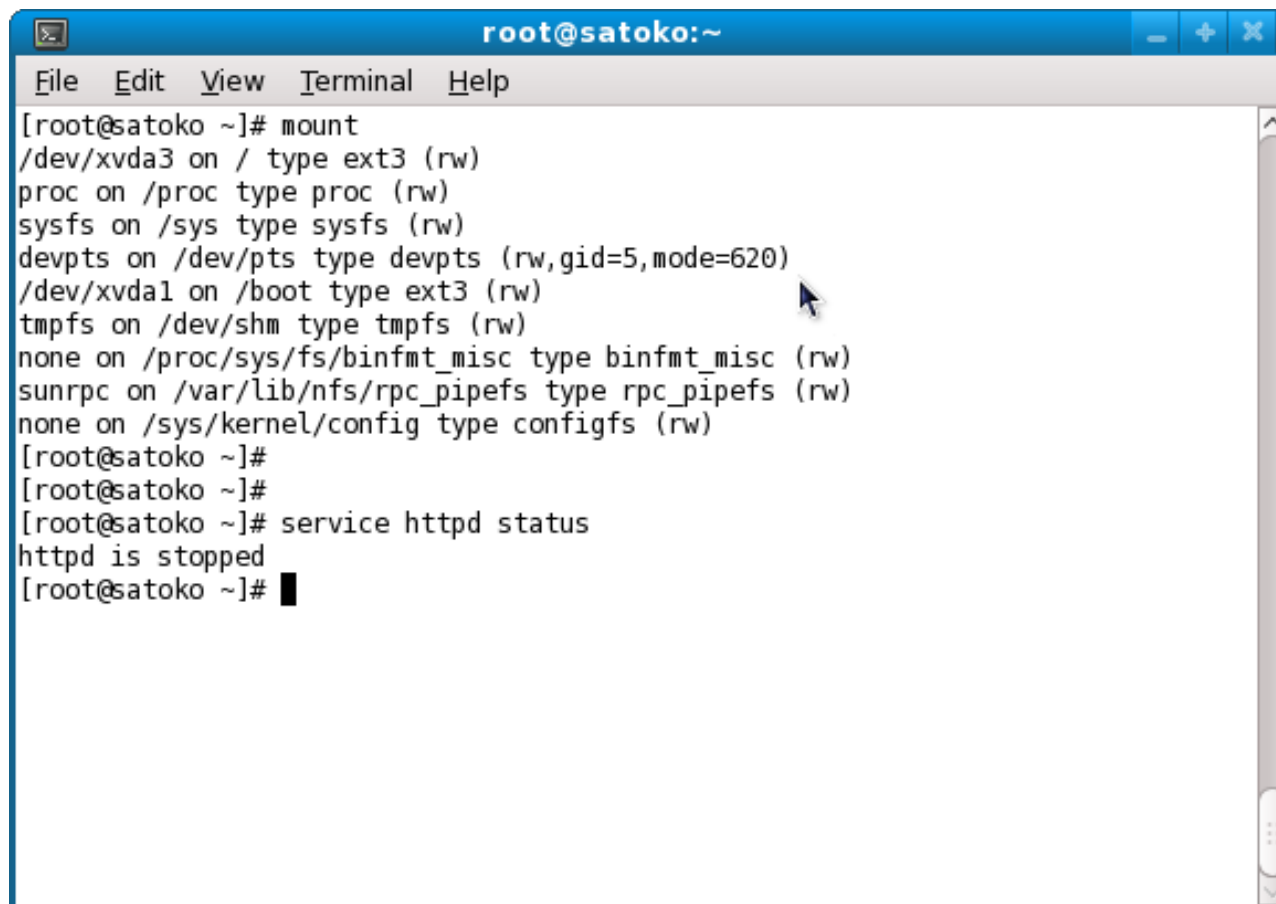
Status: Running on sakuko.salab.dfw.redhat.com
Autostart is enabled for this service

Failover Domain Association: prefer_satoko

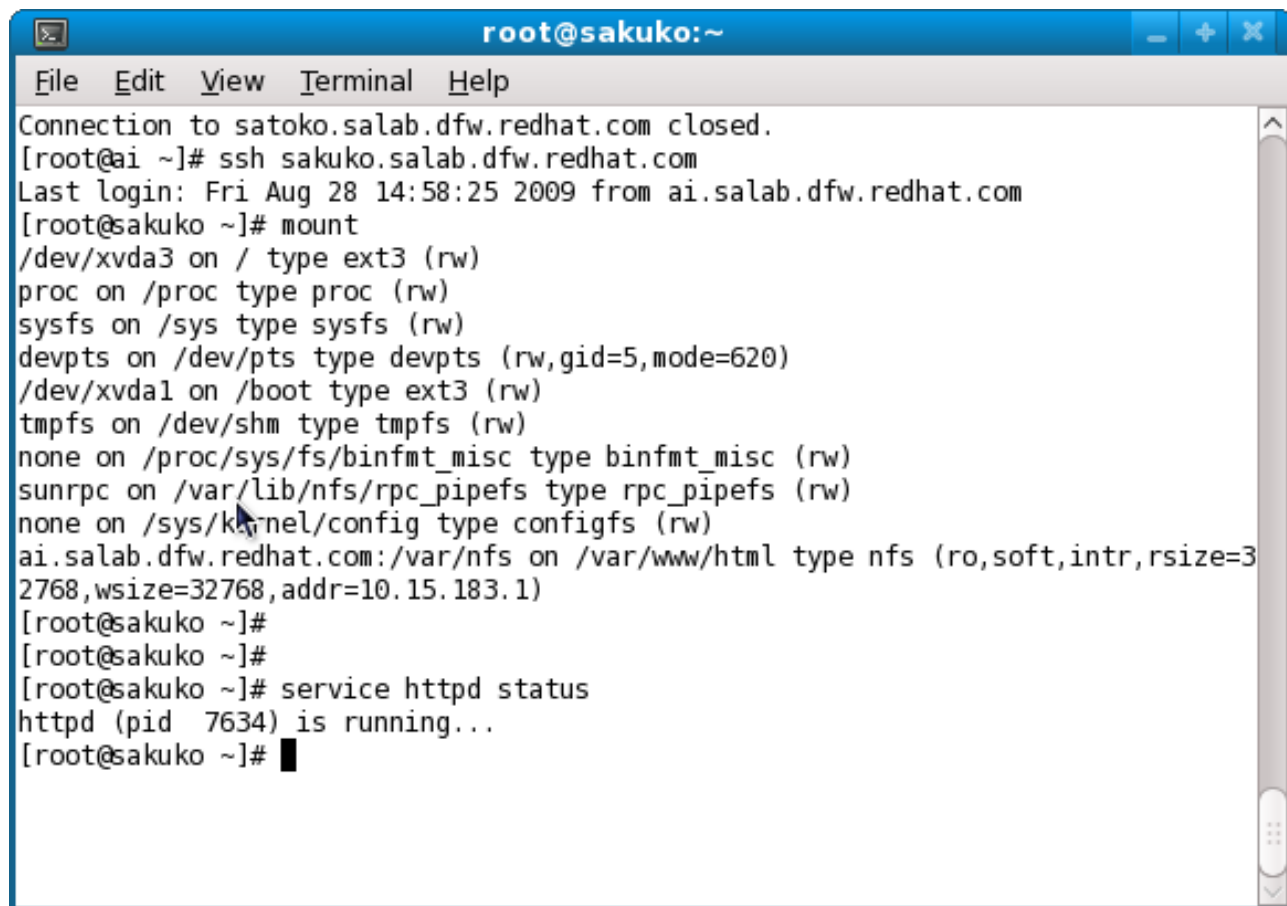
Choose a Task... Go

Find: workaro Previous Next Highlight all Match case

Done

A terminal window titled 'root@satoko:~' with a menu bar containing 'File', 'Edit', 'View', 'Terminal', and 'Help'. The terminal shows the output of the 'mount' command, listing various filesystems mounted on the system. It then shows the command 'service httpd status' and its output 'httpd is stopped'.

```
root@satoko:~  
File Edit View Terminal Help  
[root@satoko ~]# mount  
/dev/xvda3 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/xvda1 on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /sys/kernel/config type configfs (rw)  
[root@satoko ~]#  
[root@satoko ~]#  
[root@satoko ~]# service httpd status  
httpd is stopped  
[root@satoko ~]#
```

A terminal window titled 'root@sakuko:~' with a menu bar (File, Edit, View, Terminal, Help). The terminal shows the output of an SSH session and a mount command. The SSH session connects to 'sakuko.salab.dfw.redhat.com' and shows the last login time. The 'mount' command lists various filesystems mounted on the system, including /dev/xvda3 on /, /proc, /sys, /dev/pts, /dev/xvda1 on /boot, /dev/shm, and an NFS mount from ai.salab.dfw.redhat.com. The terminal ends with the command 'service httpd status' showing 'httpd (pid 7634) is running...' and a prompt for the next command.

```
root@sakuko:~  
File Edit View Terminal Help  
Connection to satoko.salab.dfw.redhat.com closed.  
[root@ai ~]# ssh sakuko.salab.dfw.redhat.com  
Last login: Fri Aug 28 14:58:25 2009 from ai.salab.dfw.redhat.com  
[root@sakuko ~]# mount  
/dev/xvda3 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/xvda1 on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /sys/kernel/config type configfs (rw)  
ai.salab.dfw.redhat.com:/var/nfs on /var/www/html type nfs (ro,soft,intr,rsize=3  
2768,wsiz=32768,addr=10.15.183.1)  
[root@sakuko ~]#  
[root@sakuko ~]#  
[root@sakuko ~]# service httpd status  
httpd (pid 7634) is running...  
[root@sakuko ~]#
```



```
root@sakuko:~  
File Edit View Terminal Help  
[root@sakuko ~]# clustat  
Cluster Status for domU-cluster @ Fri Aug 28 16:12:17 2009  
Member Status: Quorate  
  
Member Name                                ID  Status  
-----  
satoko.salab.dfw.redhat.com                1 Online, rgmanager  
sakuko.salab.dfw.redhat.com                2 Online, Local, rgmanager  
sakura.salab.dfw.redhat.com                3 Online, rgmanager  
sayoko.salab.dfw.redhat.com                4 Online, rgmanager  
sanako.salab.dfw.redhat.com                5 Online, rgmanager  
sakurako.salab.dfw.redhat.com              6 Online, rgmanager  
  
Service Name                                Owner (Last)                                State  
-----  
service:web-service                        sakuko.salab.dfw.redhat.com                started  
[root@sakuko ~]#
```

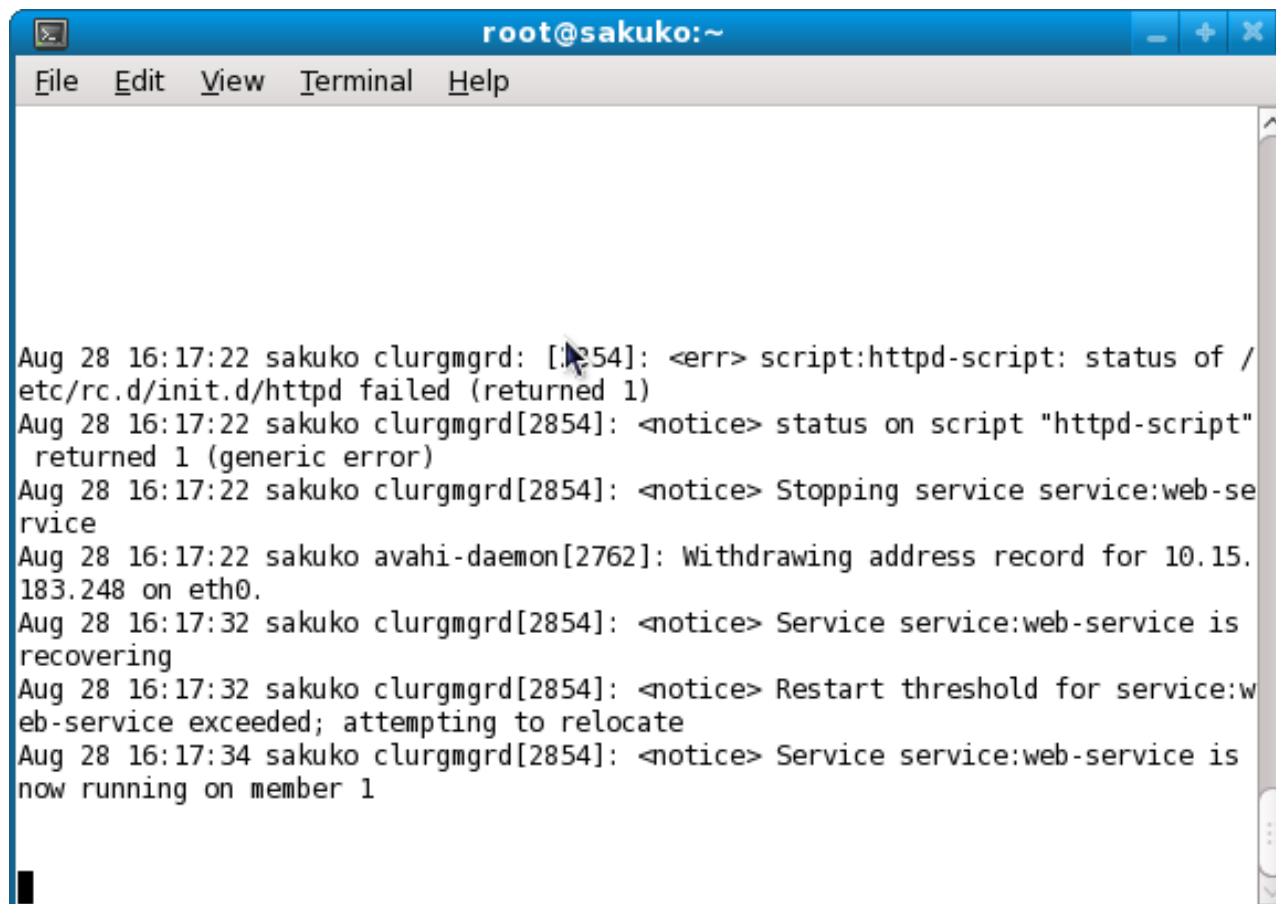
Set up the clustered app on the domU cluster

Test, test, test

Kill the app

Crash the domU

Crash the dom0



A terminal window titled "root@sakuko:~" with a menu bar containing "File", "Edit", "View", "Terminal", and "Help". The terminal displays several log messages from the system log. The messages indicate a failure of the httpd service, followed by an attempt to stop it, and then a recovery process where the service is restarted and successfully running on member 1.

```
Aug 28 16:17:22 sakuko clurgmgrd[2854]: <err> script:httpd-script: status of /  
etc/rc.d/init.d/httpd failed (returned 1)  
Aug 28 16:17:22 sakuko clurgmgrd[2854]: <notice> status on script "httpd-script"  
returned 1 (generic error)  
Aug 28 16:17:22 sakuko clurgmgrd[2854]: <notice> Stopping service service:web-se  
rvice  
Aug 28 16:17:22 sakuko avahi-daemon[2762]: Withdrawing address record for 10.15.  
183.248 on eth0.  
Aug 28 16:17:32 sakuko clurgmgrd[2854]: <notice> Service service:web-service is  
recovering  
Aug 28 16:17:32 sakuko clurgmgrd[2854]: <notice> Restart threshold for service:w  
eb-service exceeded; attempting to relocate  
Aug 28 16:17:34 sakuko clurgmgrd[2854]: <notice> Service service:web-service is  
now running on member 1
```

WOW! You made it this far!

This should give you a good start - it is not, and can not be, comprehensive.

To learn a lot more about virtualization, take RH401:

<http://tinyurl.com/5gnj2f>

To learn a lot more about clustering, look at RH436:

<http://tinyurl.com/o89lnr>

If you need help getting this all set up, contact Red Hat Consulting:

<http://www.redhat.com/consulting/>

Thank you very much!

Feel free to reach out to me at thomas@redhat.com, or Lon
at lh@redhat.com

Updated slides at

<http://people.redhat.com/tcameron>

<http://www.redhat.com/promo/summit>

Please take the survey at:

<http://www.redhat.com/summit-survey>

QUESTIONS?

**TELL US WHAT YOU THINK:
[REDHAT.COM/SUMMIT-SURVEY](https://redhat.com/summit-survey)**

RED HAT :: CHICAGO :: 2009

SUMMIT

FOLLOW US:

[TWITTER.COM/REDHATSUMMIT](https://twitter.com/redhatsummit)

TWEET ABOUT US:

ADD #SUMMIT AND/OR #JBOSSWORLD TO THE END
OF YOUR EVENT-RELATED TWEET

presented by



RED HAT :: CHICAGO :: 2009

SUMMIT

Clustering Virtual Machines and Applications with Red Hat Enterprise Linux

Thomas Cameron, RHCE, RHCX, CNE, MCSE, MCT
Solutions Architect Team Lead, Red Hat
Lon Hohberger
Senior Cluster Software Engineer, Red Hat

September 2, 2009

presented by



This Presentation

The updated version of this presentation is available at
<http://people.redhat.com/tcameron>

If you like this presentation, please take the survey at
<http://www.redhat.com/summit-survey> and tell your friends!

Agenda

Red Hat and Virtualization

Architecture

Configure Raw Storage and iSCSI

- On the Target

- On the Client

Define dom0 Cluster

- Conga

Configure Shared Storage for dom0 Machines

- Conga

Install domU machines

Agenda

Set up cluster of domU machines

Set up Application Cluster on domU machines

Questions

Red Hat and Virtualization

Red Hat leads the way in Open Source virtualization

Highly successful entry into virtualization with Xen at the release of RHEL 5.0

Acquired Qumranet, the company which developed KVM
September 2008

KVM included in RHEL 5.4 – more choice, more flexibility!

Announcing Red Hat Enterprise Virtualization (RHEV), a
KVM based hypervisor and enterprise virtualization
management system at Summit this year.

Architecture

ai - iSCSI target and Conga server. Not part of the cluster, just storage and management. Kickstarted with cluster-storage and clustering groups and scsi-target-utils package

ami, ayame and botan - iSCSI initiators, dom0 cluster members. Kickstarted with virtualization, cluster-storage and clustering groups and iscsi-initiator-utils package


sakuko, sakura, sakurako, sanako, satoko and sayoko - domU machines. Kickstarted with cluster-storage and clustering groups. Running an Apache cluster.

All systems are on a gigabit ethernet network

```
install
text
network --bootproto dhcp
url --url http://cho.salab.dfw.redhat.com/ty/xZSxYSae
lang en_US
langsupport --default en_US en_US
keyboard us
mouse none
zerombr yes
clearpart --all
part /boot --fstype=ext3 --size=100
part pv.01 --size=1000 --grow
part swap --size=2100
volgroup VolGroup00 pv.01
logvol / --vgname=VolGroup00 --name=LogVol00 --size=1 --grow
bootloader --location mbr
timezone --utc America/Chicago
auth --enablemd5 --enablesshadow
rootpw --iscrypted $1$NSMQB/$4hPK9.OkxSs/bx3al.l72/
selinux --disabled
reboot
firewall --disabled
skipx
repo --name=Cluster --baseurl=http://cho.salab.dfw.redhat.com/kickstart/dist/ks-rhel-x86_64-server-5-u3/Cluster
repo --name=ClusterStorage --baseurl=http://cho.salab.dfw.redhat.com/kickstart/dist/ks-rhel-x86_64-server-5-u3/ClusterStorage
repo --name=VT --baseurl=http://cho.salab.dfw.redhat.com/kickstart/dist/ks-rhel-x86_64-server-5-u3/VT
repo --name=Workstation --baseurl=http://cho.salab.dfw.redhat.com/kickstart/dist/ks-rhel-x86_64-server-5-u3/Workstation
key --skip

%packages --resolvedeps

@virtualization
kernel-xen
@clustering
@cluster-storage
@base
xen
scsi-target-utils
```



Configure storage and iSCSI on ai

Create the physical volume, volume group and logical volume. In this case, /dev/sdb1 is a newly created, blank partition of 1TiB:

```
pvcreate /dev/sdb1
vgcreate VolGroup01 /dev/sdb1
lvcreate -n VMStorage -L 100GB VolGroup01
```

/dev/sdb1 is the whole disk.

We create one large physical volume with the pvcreate command.

We use the whole physical volume (and therefore the whole disk) for the volume group with the vgcreate command.

Then we create a small, 100GiB logical volume called VMStorage with the lvcreate command.

Configure storage and iSCSI on ai

Make sure that the iSCSI target software is installed

```
yum -y groupinstall clustering cluster-storage
```

```
yum -y install scsi-target-utils
```

```
chkconfig tgtd on
```

Could easily get away with only installing the luci RPM instead of all of the clustering and cluster-storage groups.

scsi-target-utils will install perl-Config-General as a dependency

Only chkconfig tgtd on, don't start it yet.

Configure storage and iSCSI on ai

Configure the target

Configure the /etc/tgtd/target.conf to look something like this:

```
<target iqn.2009-06.com.redhat.dfw.salab:disk1>  
    backing-store /dev/mapper/VolGroup01-VMStorage  
    initiator-address 10.15.183.6  
    initiator-address 10.15.183.7  
    initiator-address 10.15.183.8  
</target>
```

The iSCSI Qualified Name (iqn) is the string “iqn.[the date].[the domain name in reverse]:[disk]”

The backing store can be any block device or logical volume.

The initiator address lines are there for security - only allow those servers to connect which are authorized to do so

Configure storage and iSCSI on ai

Stop and start the tgtd service on the iSCSI target

```
service tgtd stop
```

```
sleep 5
```

```
service tgtd start
```

I usually do a 5 second sleep between stopping and starting, otherwise there is a potential for the target to not log out cleanly and then fail to start

Configure storage and iSCSI on ai

Check that your target is available

```
tgtadm --lld iscsi --op show --mode target
```

Will give you a verbose readout of the targets available

```
root@ai:~  
File Edit View Terminal Help  
[root@ai ~]# tgtadm --lld iscsi --op show --mode target  
Target 1: iqn.2009-06.com.redhat:dfw.slab:disk1  
System information:  
  Driver: iscsi  
  State: ready  
I_T nexus information:  
  I_T nexus: 26  
    Initiator: iqn.1994-05.com.redhat:88647efd3573  
    Connection: 0  
    IP Address: 10.15.183.6  
  I_T nexus: 27  
    Initiator: iqn.1994-05.com.redhat:d6eff839a065  
    Connection: 0  
    IP Address: 10.15.183.8  
  I_T nexus: 28  
    Initiator: iqn.1994-05.com.redhat:437eb150453a  
    Connection: 0  
    IP Address: 10.15.183.7  
LUN information:  
  LUN: 0  
    Type: controller  
    SCSI ID: deadbeaf1:0  
    SCSI SN: beaf10  
    Size: 0 MB  
    Online: Yes  
    Removable media: No  
    Backing store: No backing store  
  LUN: 1  
    Type: disk  
    SCSI ID: deadbeaf1:1  
    SCSI SN: beaf11  
    Size: 107374 MB  
    Online: Yes  
    Removable media: No  
    Backing store: /dev/mapper/VolGroup01-VMStorage  
Account information:  
ACL information:  
  10.15.183.6  
  10.15.183.7  
  10.15.183.8  
[root@ai ~]#
```

Configure storage and iSCSI on initiators

On each of the dom0 servers (ami, ayame and botan)

```
yum -y groupinstall clustering cluster-storage  
yum -y install iscsi-initiator-utils  
chkconfig iscsi on  
chkconfig iscsid on  
service iscsid restart  
service iscsi restart
```

In my lab, I kickstarted ami, ayame and botan with the @ clustering and @ cluster-storage groups and iscsi-initiator-utils, so the first two steps are actually redundant

Make sure you (re)start iscsid before iscsi. iscsid inserts several necessary kernel modules and checks network setting sanity. The iscsi service only handles logins.

Configure storage and iSCSI on initiators

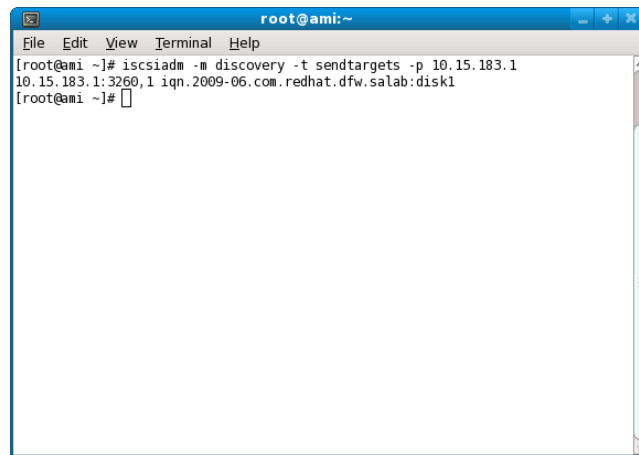
Discover the iSCSI target

```
iscsiadm -m discovery -t sendtargets -p 10.15.183.1
```

-m discovery sets the iscsi command mode to discover

-t sendtargets asks the target to send iSCSI target info to the initiator

-p 10.15.183.1 is the address of the portal. Optionally you can add a port number with the syntax “-p 10.15.183.1:3260” 3260 is the default port.



A terminal window titled 'root@ami:~' with a menu bar (File, Edit, View, Terminal, Help). The terminal shows the command `iscsiadm -m discovery -t sendtargets -p 10.15.183.1` being executed, followed by the output `10.15.183.1:3260,1 iqn.2009-06.com.redhat.dfw.salab:disk1`. The prompt returns to `root@ami ~]#`.

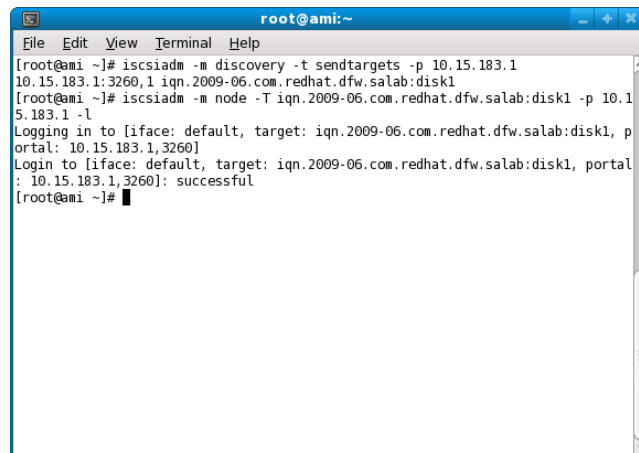
```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# iscsiadm -m discovery -t sendtargets -p 10.15.183.1  
10.15.183.1:3260,1 iqn.2009-06.com.redhat.dfw.salab:disk1  
[root@ami ~]#
```

Configure storage and iSCSI on initiators

Log in to the target (command below is actually all one command, it is on multiple lines for ease of visibility)

```
iscsiadm -m node \  
-T iqn.2009-06.com.redhat.dfw.salab:disk1 \  
-p 10.15.183.1 -l
```

- m node sets the mode to that of a node on the cluster
- T [iqn] specifies the target we got from the previous command
- p [address] is the address and, optionally, the port number of the portal (target)
- l logs on to the target

A terminal window titled 'root@ami:~' with a menu bar (File, Edit, View, Terminal, Help). The terminal shows the following commands and output:

```
[root@ami ~]# iscsiadm -m discovery -t sendtargets -p 10.15.183.1
10.15.183.1:3260,1 iqn.2009-06.com.redhat.dfw.salab:disk1
[root@ami ~]# iscsiadm -m node -T iqn.2009-06.com.redhat.dfw.salab:disk1 -p 10.15.183.1 -l
Logging in to [iface: default, target: iqn.2009-06.com.redhat.dfw.salab:disk1, portal: 10.15.183.1,3260]
Login to [iface: default, target: iqn.2009-06.com.redhat.dfw.salab:disk1, portal: 10.15.183.1,3260]: successful
[root@ami ~]#
```

In this slide we took the output of the previous discovery and used it to do a login

- m node sets the mode of the action to be that of a node on the storage cluster

- T [iqn] is the iqn of the LUN we want to attach to

- p [address] is the address of the target

You should see the “Login to iface...” statement end in “successful”

Configure storage and iSCSI on initiators

Run “fdisk -l” or “cat /proc/partitions” to verify that your system “sees” the new disk

There should be no partition table on e.g. /dev/sdb

```
root@ami:~  
File Edit View Terminal Help  
10.15.183.1:3260,1 iqn.2009-06.com.redhat.dfw.salab:disk1  
[root@ami ~]# iscsiadm -m node -T iqn.2009-06.com.redhat.dfw.salab:disk1 -p 10.15.183.1 -l  
Logging in to [iface: default, target: iqn.2009-06.com.redhat.dfw.salab:disk1, portal: 10.15.183.1,3260]  
Login to [iface: default, target: iqn.2009-06.com.redhat.dfw.salab:disk1, portal: 10.15.183.1,3260]: successful  
[root@ami ~]# fdisk -l  
  
Disk /dev/sda: 1000.2 GB, 1000204886016 bytes  
255 heads, 63 sectors/track, 121601 cylinders  
Units = cylinders of 16065 * 512 = 8225280 bytes  


| Device    | Boot | Start | End    | Blocks    | Id | System               |
|-----------|------|-------|--------|-----------|----|----------------------|
| /dev/sda1 | *    | 1     | 13     | 104391    | 83 | Linux                |
| /dev/sda2 |      | 14    | 281    | 2152710   | 82 | Linux swap / Solaris |
| /dev/sda3 |      | 282   | 121601 | 974502900 | 8e | Linux LVM            |

  
Disk /dev/sdb: 107.3 GB, 107374182400 bytes  
255 heads, 63 sectors/track, 13054 cylinders  
Units = cylinders of 16065 * 512 = 8225280 bytes  
  
Disk /dev/sdb doesn't contain a valid partition table  
[root@ami ~]#
```

Configure storage and iSCSI on initiators

Create a partition on the LUN using fdisk

```
fdisk /dev/sdb
```

Choose n for new

Choose p for primary

Choose 1 for first partition

Hit [enter] to start at the first cylinder

Hit [enter] to end at the last cylinder

Hit w to write changes to disk

As root, run “partprobe” and “fdisk -l” on the other nodes and verify that the new partition shows up.

This example assumes use of the whole LUN.

```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# fdisk /dev/sdb  
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel  
Building a new DOS disklabel. Changes will remain in memory only,  
until you decide to write them. After that, of course, the previous  
content won't be recoverable.  
  
The number of cylinders for this disk is set to 13054.  
There is nothing wrong with that, but this is larger than 1024,  
and could in certain setups cause problems with:  
1) software that runs at boot time (e.g., old versions of LILO)  
2) booting and partitioning software from other OSs  
   (e.g., DOS FDISK, OS/2 FDISK)  
Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)  
  
Command (m for help): n  
Command action  
   e   extended  
   p   primary partition (1-4)  
p  
Partition number (1-4): 1  
First cylinder (1-13054, default 1):  
Using default value 1  
Last cylinder or +size or +sizeM or +sizeK (1-13054, default 13054):  
Using default value 13054  
Command (m for help):
```

Configure storage and iSCSI on initiators

Make sure that the ricci service is installed and running on all the dom0 nodes.

```
service ricci status
```

Again, on my lab system, since they were built with the clustering and cluster-storage groups, ricci was running at boot time.

Configure the Cluster Using Conga

On the management station (ai), make sure that the luci package is installed

```
yum -y install luci
```

```
chkconfig luci on
```

Don't start the luci service yet!

Add an administrator password

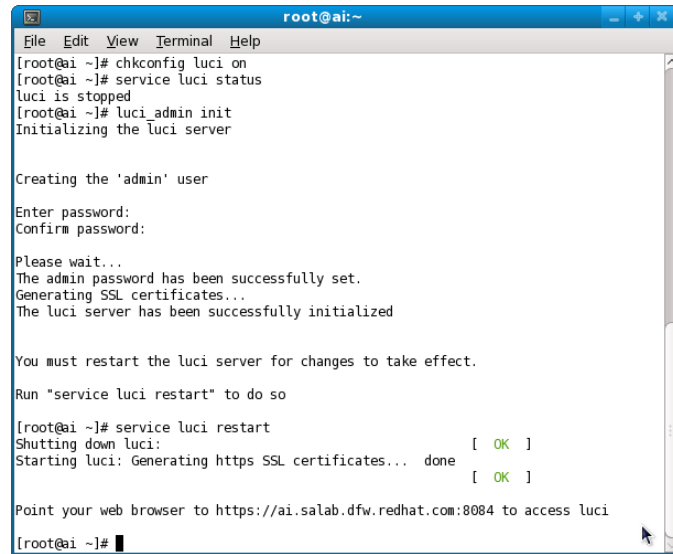
```
luci_admin init
```

```
service luci start
```

Since my systems were kickstarted with the clustering and cluster-storage package groups installed, luci was already there.

Make sure that luci is not started when running "luci_admin init"

Go to the URL listed by the restart to access Conga

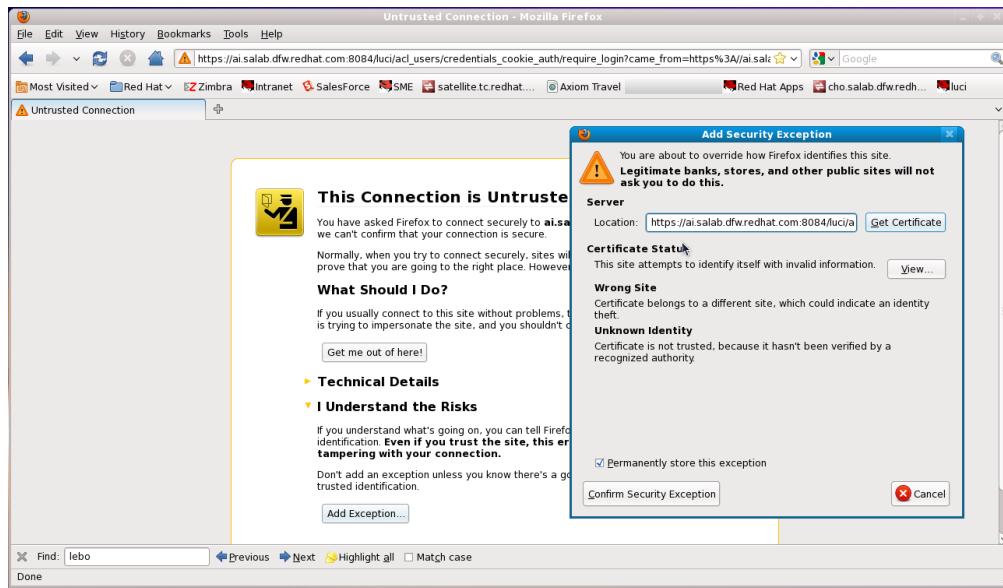


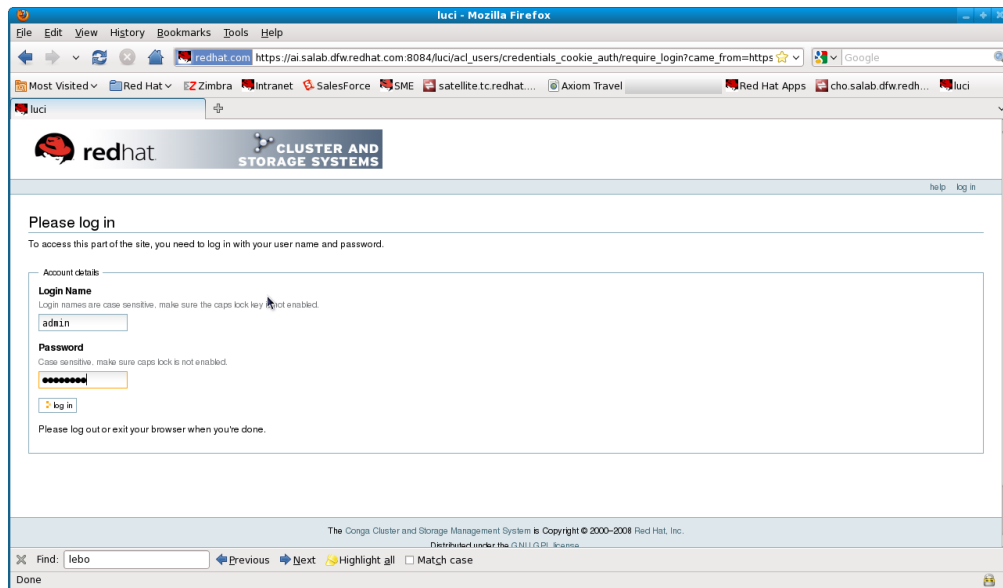
```
root@ai:~  
File Edit View Terminal Help  
[root@ai ~]# chkconfig luci on  
[root@ai ~]# service luci status  
luci is stopped  
[root@ai ~]# luci_admin init  
Initializing the luci server  
  
Creating the 'admin' user  
Enter password:  
Confirm password:  
  
Please wait...  
The admin password has been successfully set.  
Generating SSL certificates...  
The luci server has been successfully initialized  
  
You must restart the luci server for changes to take effect.  
Run "service luci restart" to do so  
  
[root@ai ~]# service luci restart  
Shutting down luci: [ OK ]  
Starting luci: Generating https SSL certificates... done [ OK ]  
  
Point your web browser to https://ai.salab.dfw.redhat.com:8084 to access luci  
[root@ai ~]#
```

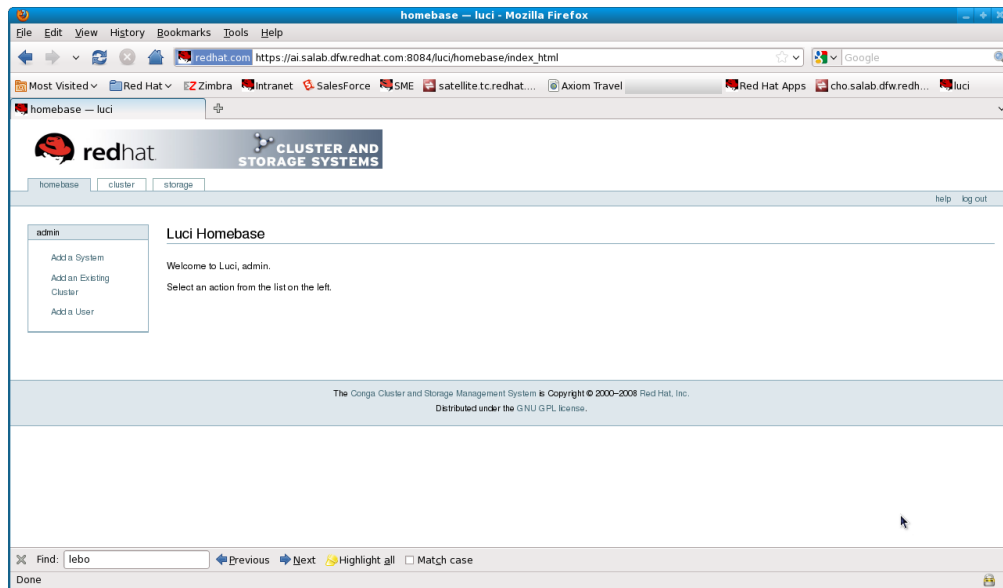

Configure the Cluster Using Conga

Open the URL listed by the “service luci start” command in your web browser

You will be warned about the self-signed certificate, it's fine to accept it assuming you trust your network







Configure the Cluster Using Conga

Choose the “Cluster” tab and create a new cluster

- Cluster name

- Node names and passwords

- Download or use locally installed

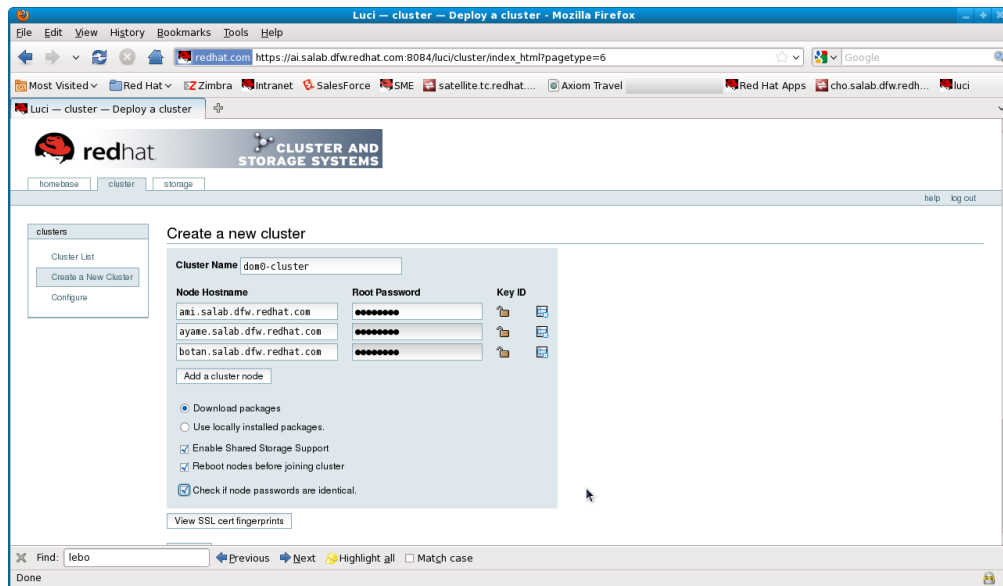
- Enable shared storage since we'll be using GFS

- Reboot nodes

- Check the box if the passwords are all the same

I always choose “download packages” just on the off chance there is a newer version on RHN

I always choose to reboot the nodes just to make sure the cluster comes up correctly

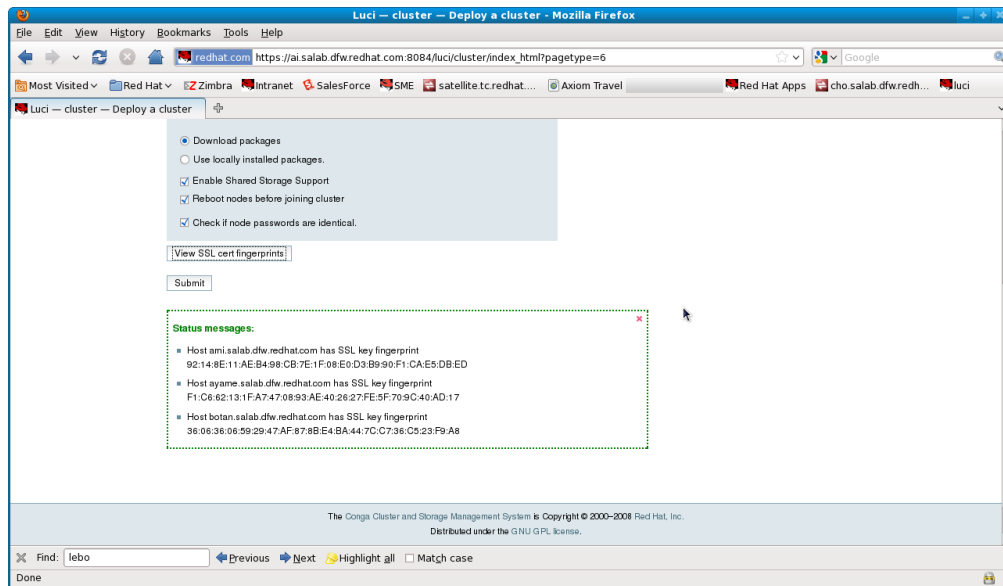


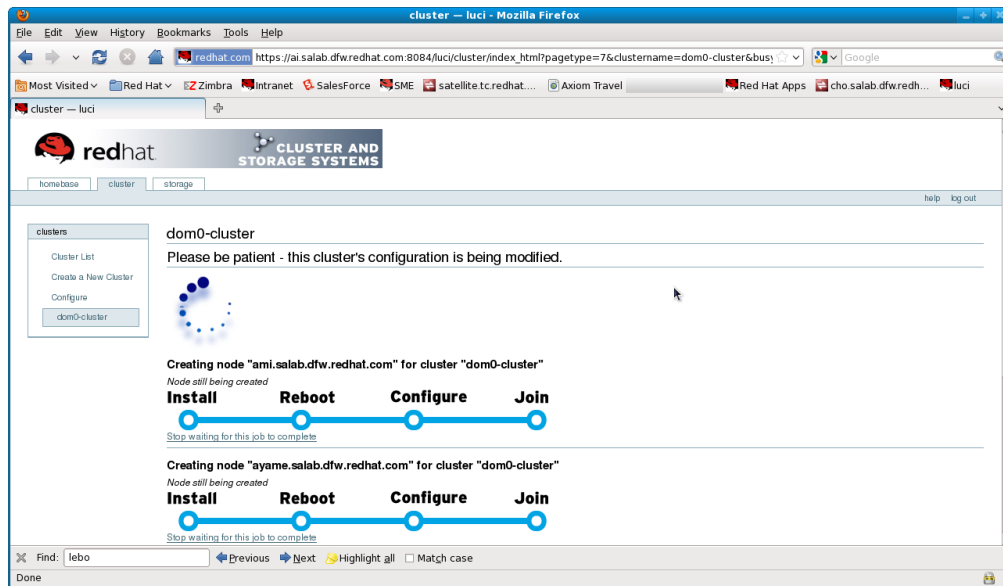
Configure the Cluster Using Conga

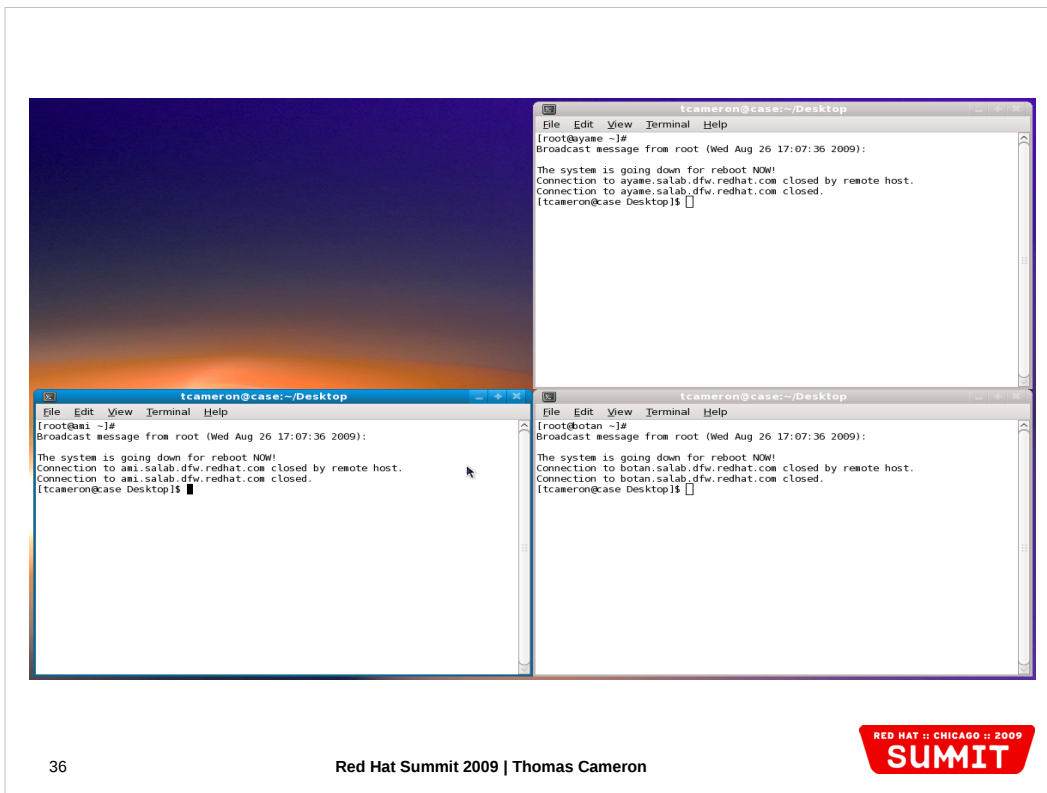
Check to see if the luci server can talk to the ricci services on all nodes

Click “View SSL fingerprints” button

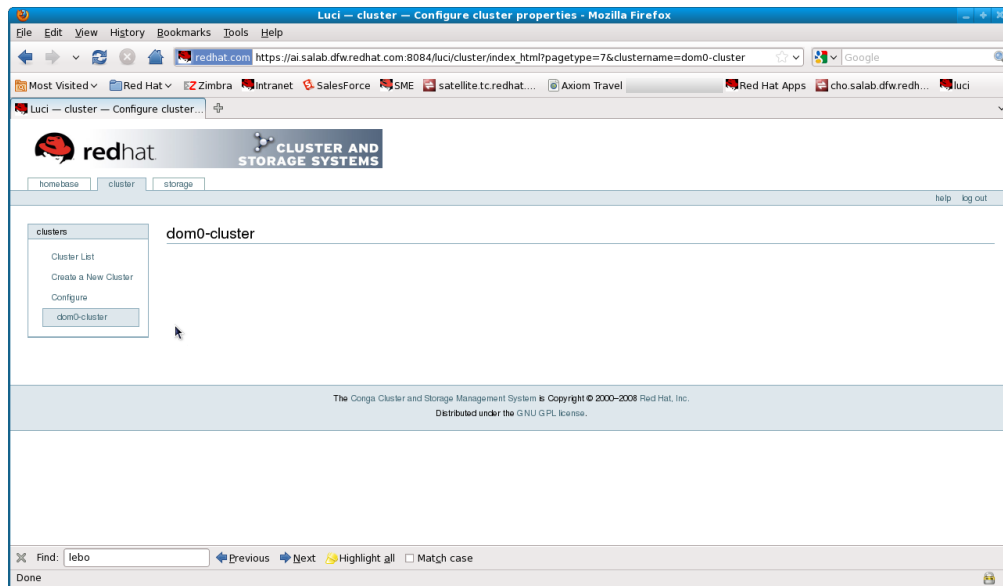
If that checks out, click on “Submit”







The three dom0 machines reboot.

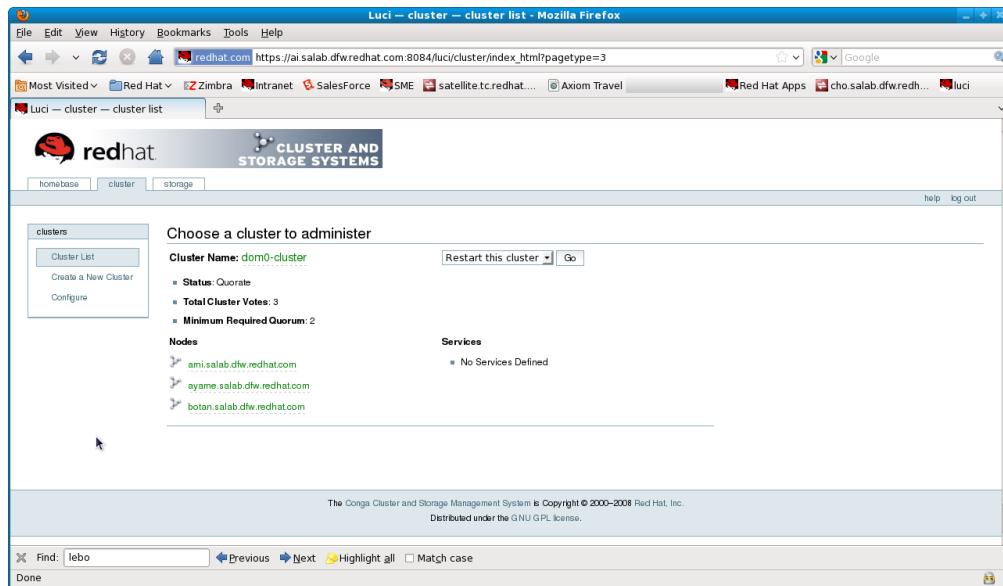


Configure the Cluster Using Conga

After the nodes all reboot, click on “Cluster List” to see if everything came up

If the cluster name and the node names are **green**, you're good.

Note that it is not unusual for the web UI to throw an error until the nodes reboot, since luci can't reach ricci on the nodes.



Configure the Cluster Using Conga

Set up a shared fence device

- Cluster tab

- Cluster name

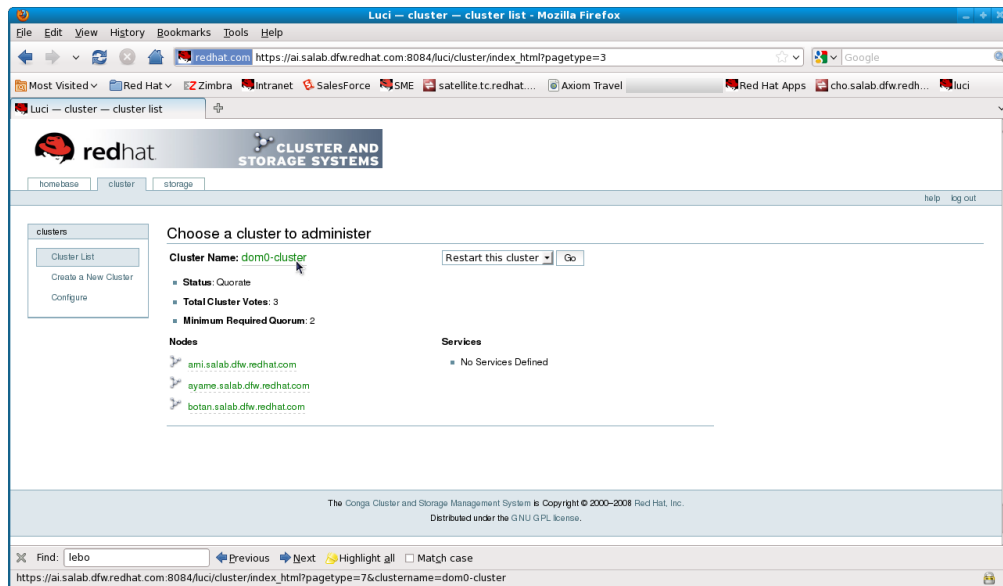
- Shared Fence Devices

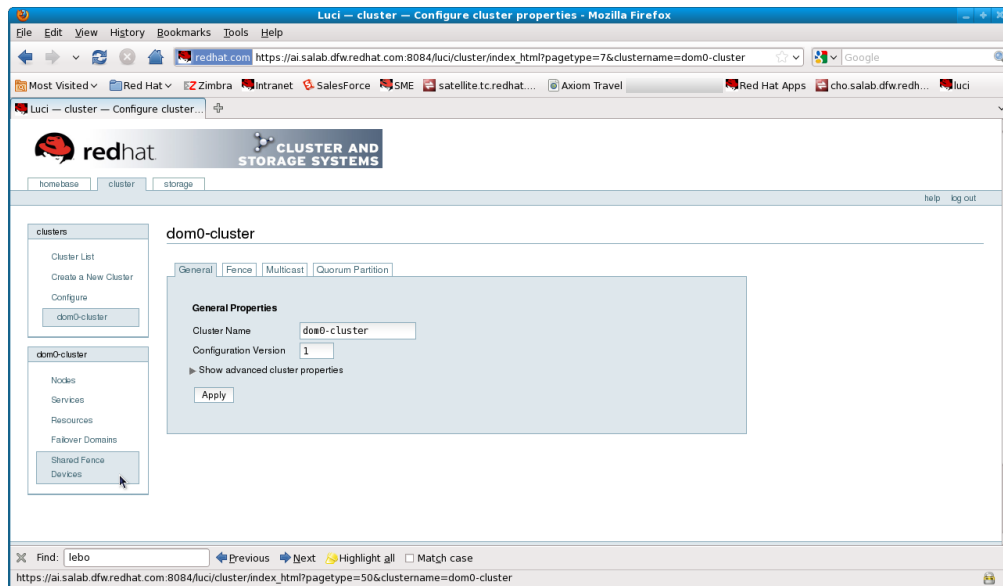
- Add a shareable fence device

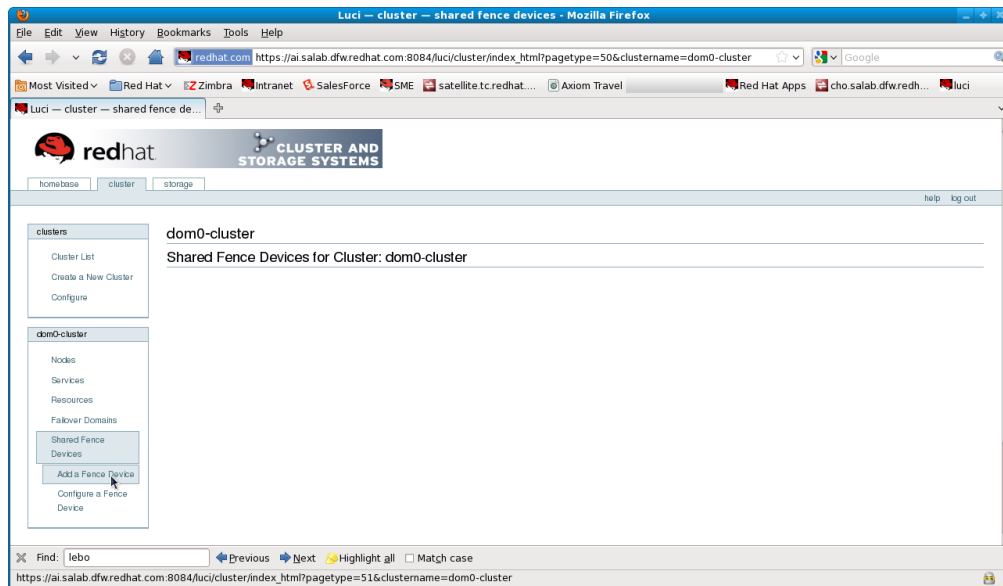
- Fill out details

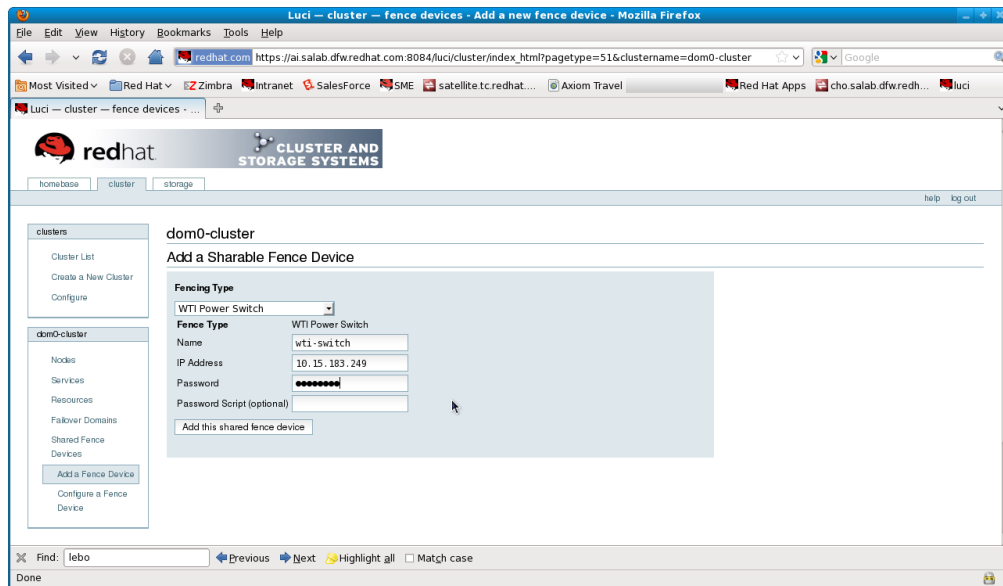
- Submit and accept

There are numerous fence devices supported, including fiber switch, remote power switch, iLO, DRAC, etc.









Configure the Cluster Using Conga

Define virtual machine fencing

Cluster tab

Click on cluster

Choose Fence tab

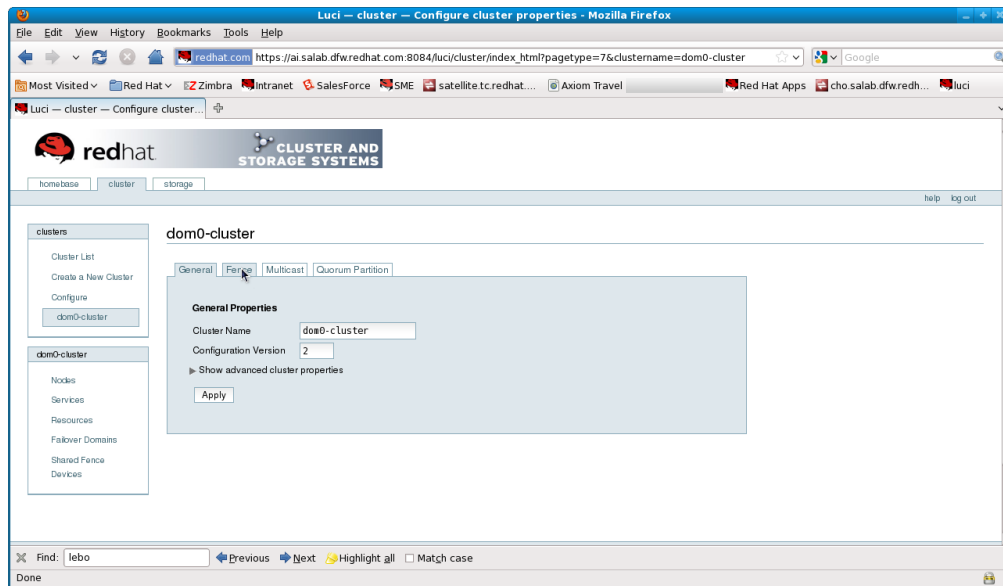
Check the box for XVM fence daemon

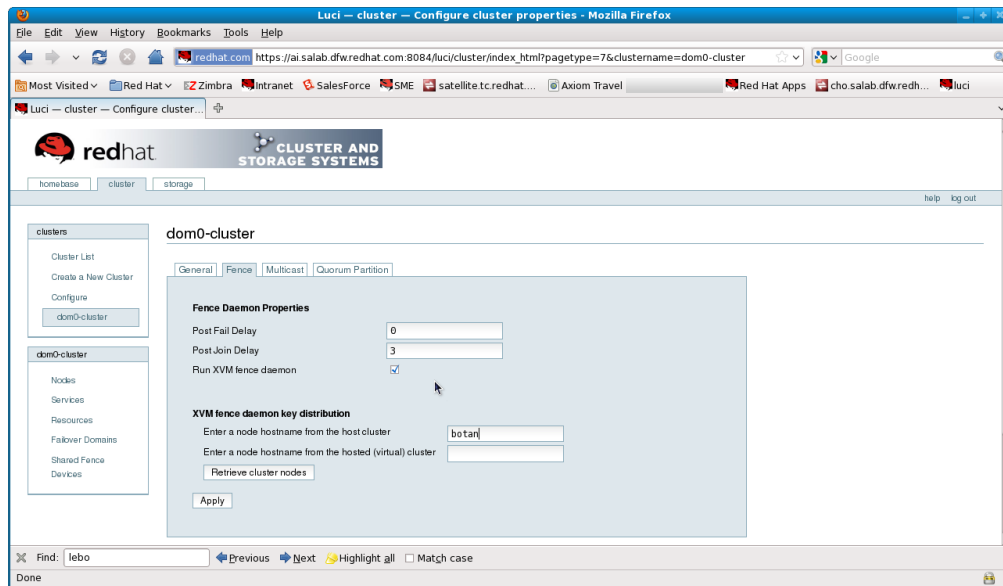
Enter a hostname from dom0 cluster

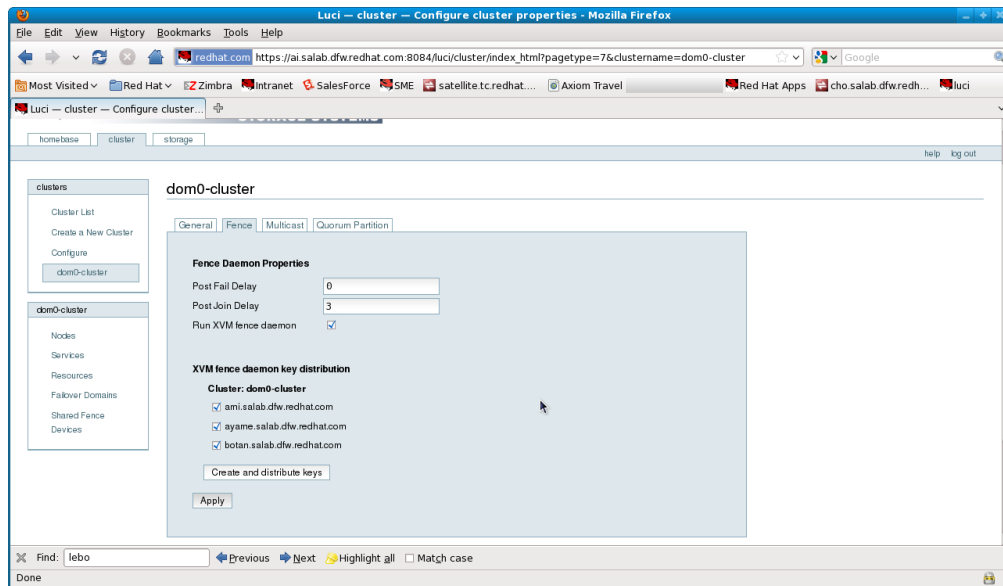
- Retrieve cluster nodes

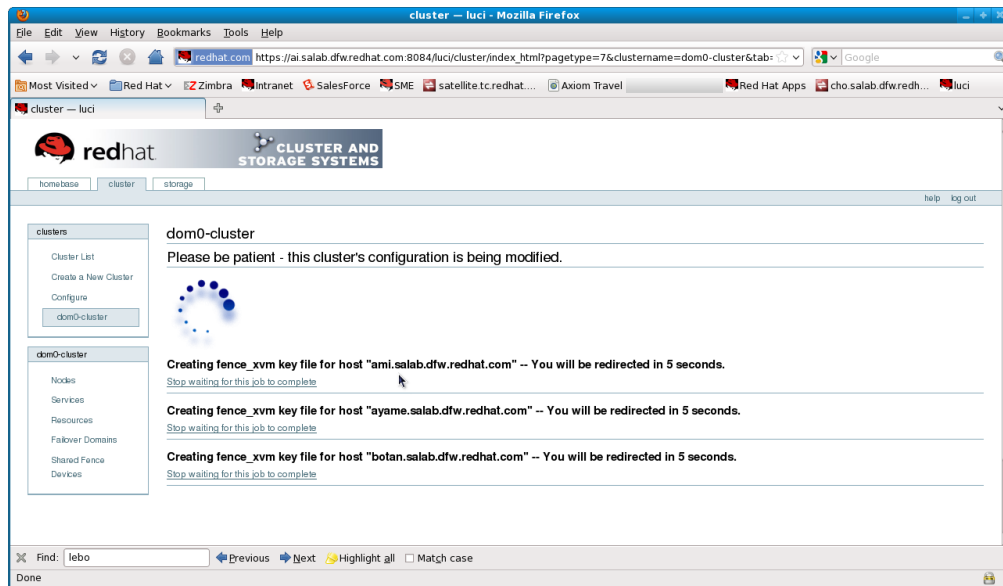
- Create and distribute keys

Apply and confirm



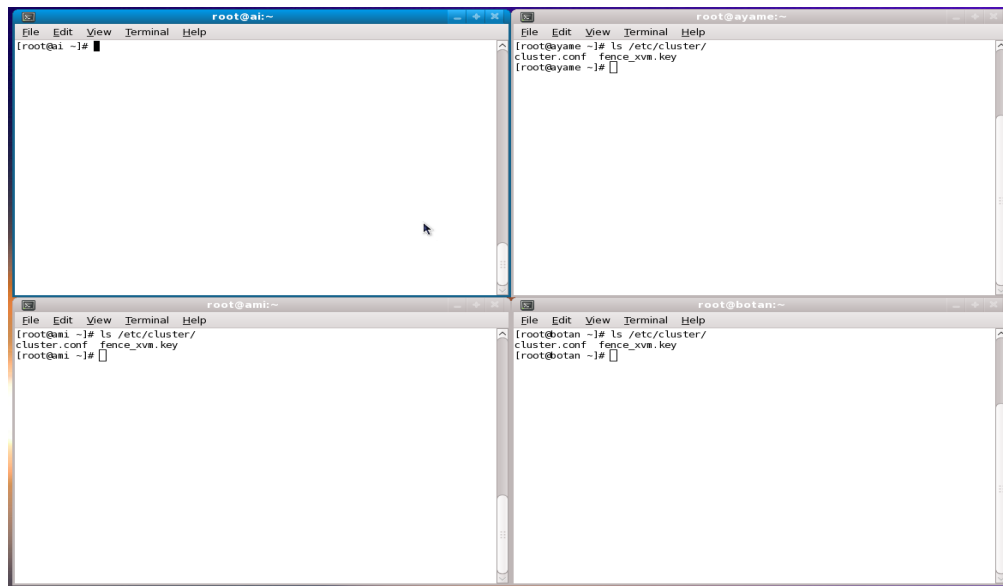






Configure the Cluster Using Conga

Check that `/etc/fence_xvm.key` was created on each of the dom0 cluster node members.



Configure Per-Node Fencing

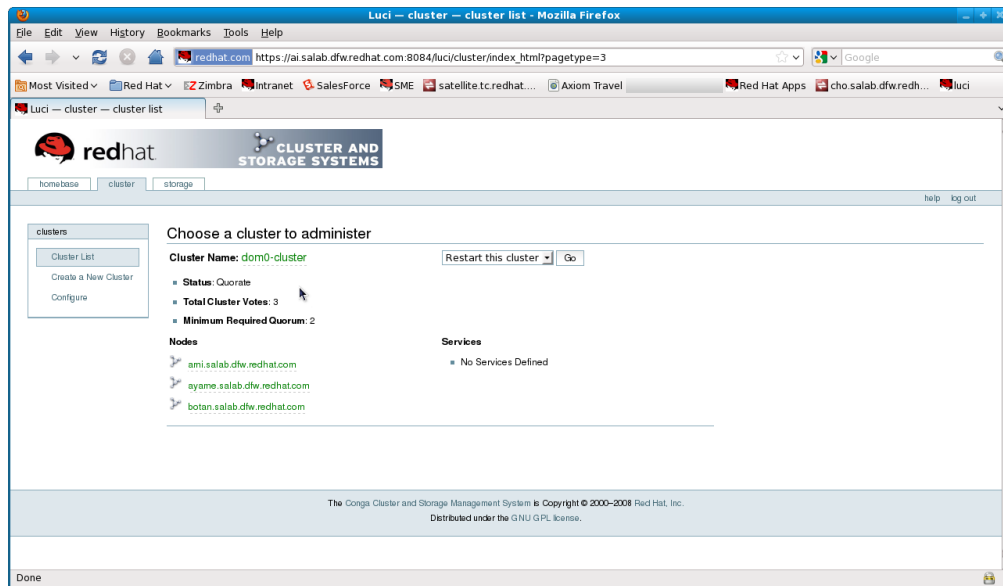
Cluster tab

Cluster Name

Node Name

Main Fencing Method

Add a fence device to this level



Luci — cluster — nodes — properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.com https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&clustername=dom0-cluster&node=...

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.t.c.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — nodes — proper...

clusters

- Cluster List
- Create a New Cluster
- Configure

dom0-cluster

Nodes

- Add a Node
- Configure
- ami.salab.dfw.redhat.com
- ayame.salab.dfw.redhat.com
- botan.salab.dfw.redhat.com

Services

- Resources
- Fallover Domains
- Shared Fence Devices

dom0-cluster

Node Name: ami.salab.dfw.redhat.com Choose a Task... Go

Status: Cluster member

Show recent log activity for this node

Cluster daemons running on this node

| Daemon | Currently running | Enabled at start-up |
|-----------|-------------------|-------------------------------------|
| cman | yes | <input checked="" type="checkbox"/> |
| rgmanager | yes | <input checked="" type="checkbox"/> |

[Update node daemon properties](#)

Services on this Node

- No cluster services are currently running here

Fallover Domain Membership

- This node has no fallover domain membership

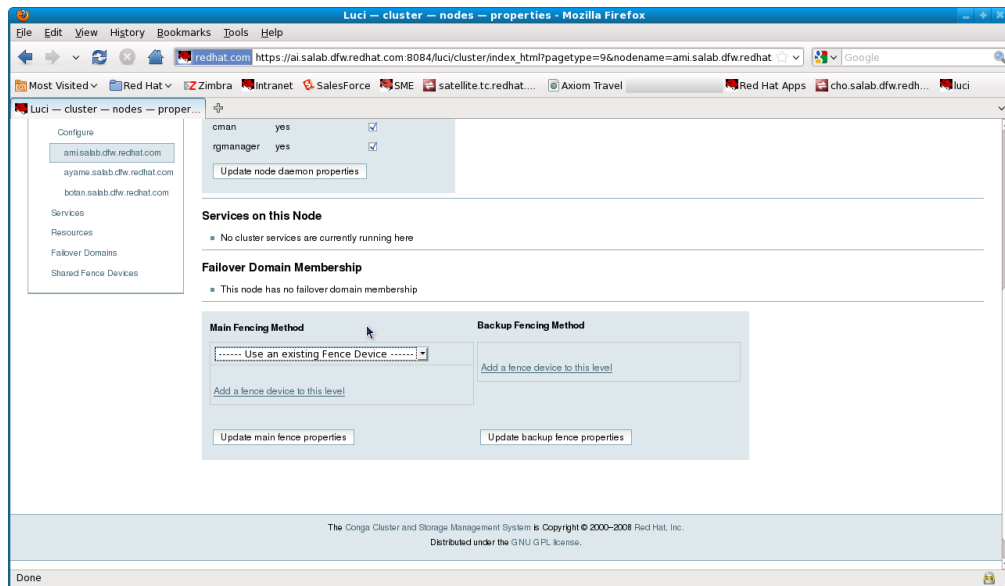
Main Fencing Method

[Add a fence device to this level](#)

Backup Fencing Method

[Add a fence device to this level](#)

Done



When you click the drop-down, choose the existing fencing device. In this case, a WTI power switch.

Luci — cluster — nodes — properties — Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.comhttps://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&nodename=ami.salab.dfw.redhat

Most VisitedRed HatZimbraIntranetSalesForceSMESatellite.t.c.redhat...Axiom TravelRed Hat Appscho.salab.dfw.redh...luci

Luci — cluster — nodes — proper...

ServicesResourcesFailover DomainsShared Fence Devices

Services on this Node

- No cluster services are currently running here

Failover Domain Membership

- This node has no failover domain membership

Fence TypeWTI Power Switch

Namewtl-swit ch

IP Address10.15.183.249

Password

Password Script (optional)

Port

Use SSH

Remove this instance

Remove this deviceAdd an instance

Add a fence device to this level

Update main fence properties

Backup Fence Method

Add a fence device to this level

Update backup fence properties

Done

56

Red Hat Summit 2009 | Thomas Cameron

RED HAT :: CHICAGO :: 2009
SUMMIT

If you've already specified the password for the power switch, and there are no per-port passwords, you do not need to set the password again. Just the power port.

Configure Per-Node Fencing

Lather, rinse, repeat for each node.

Make sure that you put the right power port number in for each node.

Configure Storage Using Conga

Set up the volume group which will be used to house the virtual machine

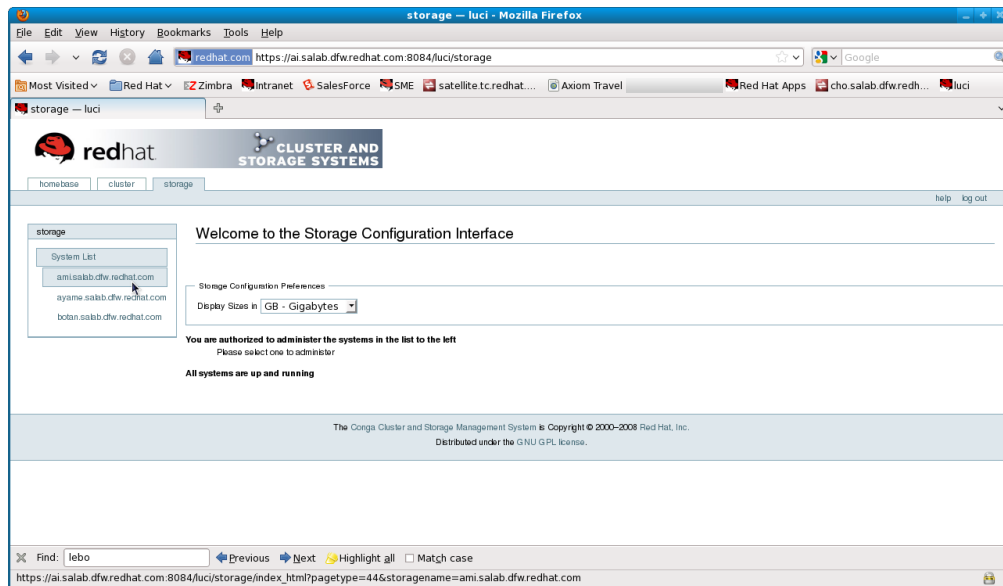
- Storage tab

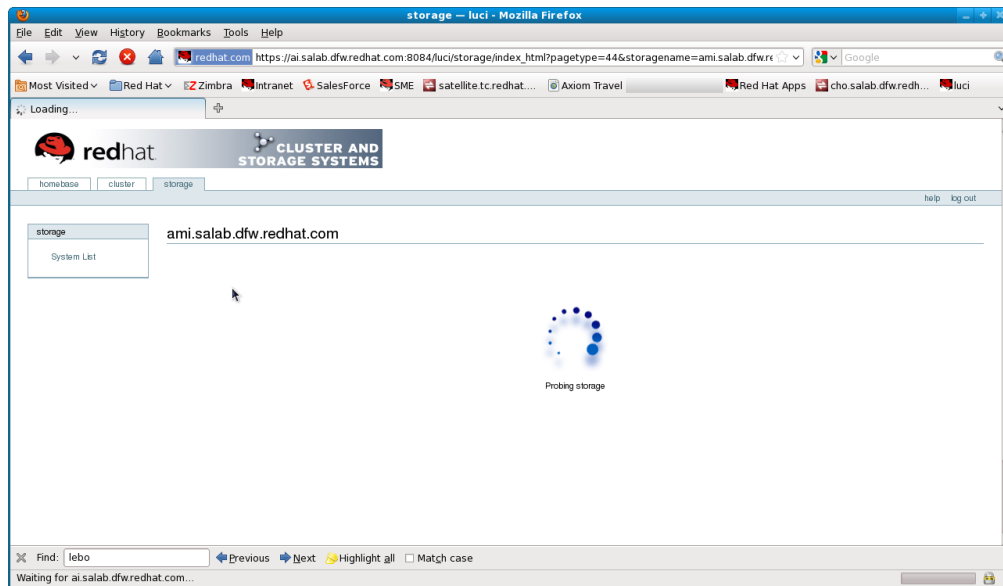
- Choose a server

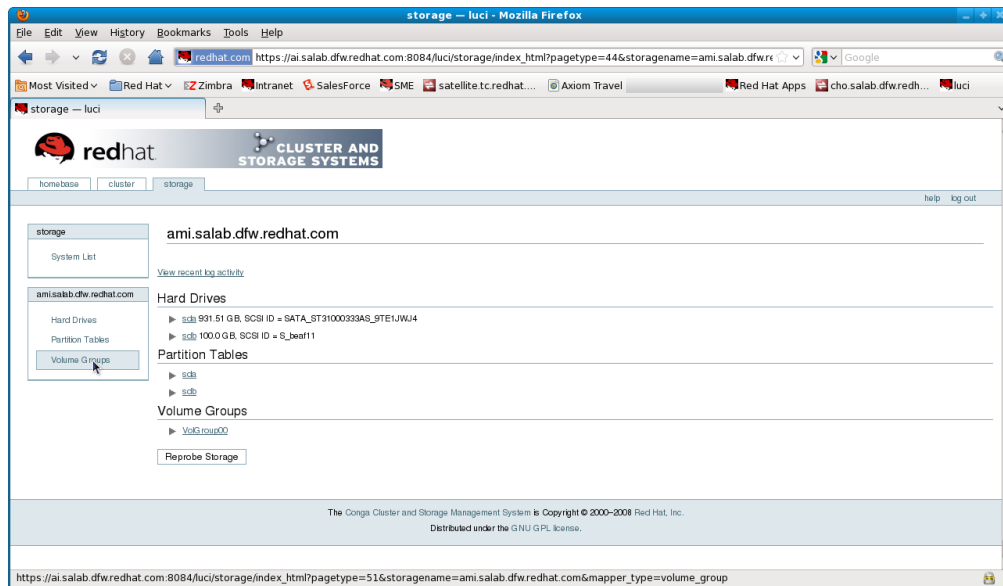
- Choose Volume Groups

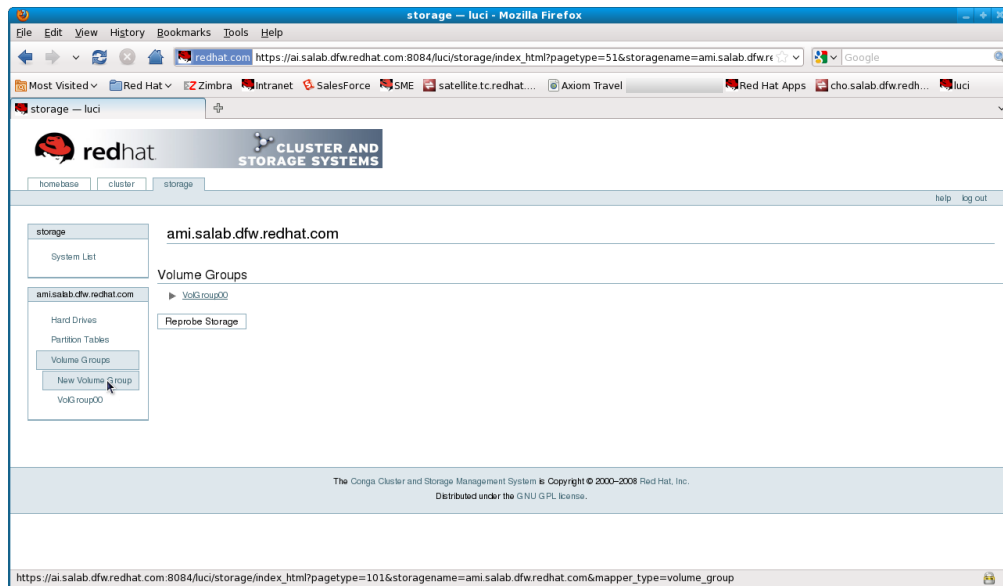
- Choose New Volume Group

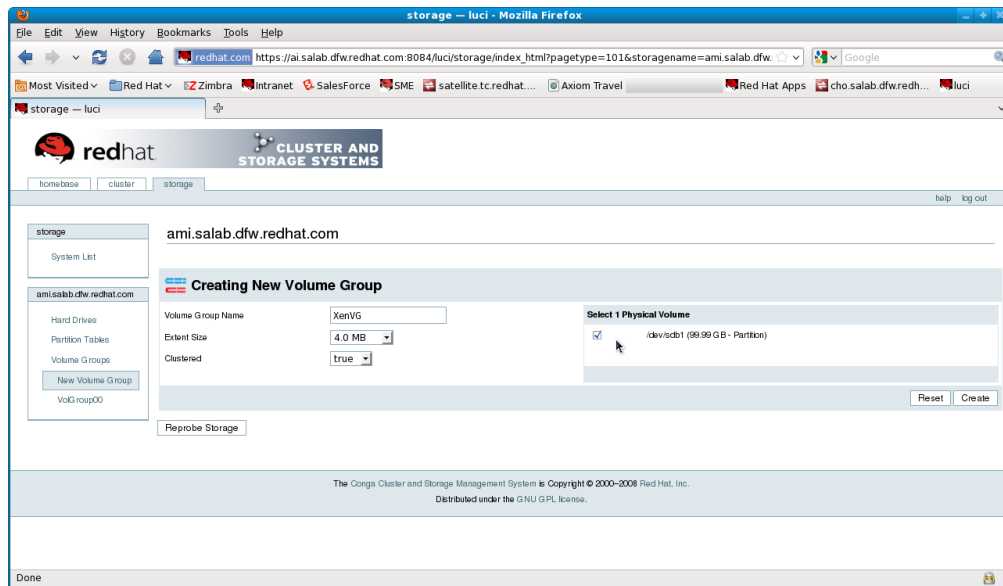
- Name new VG and assign to a block device

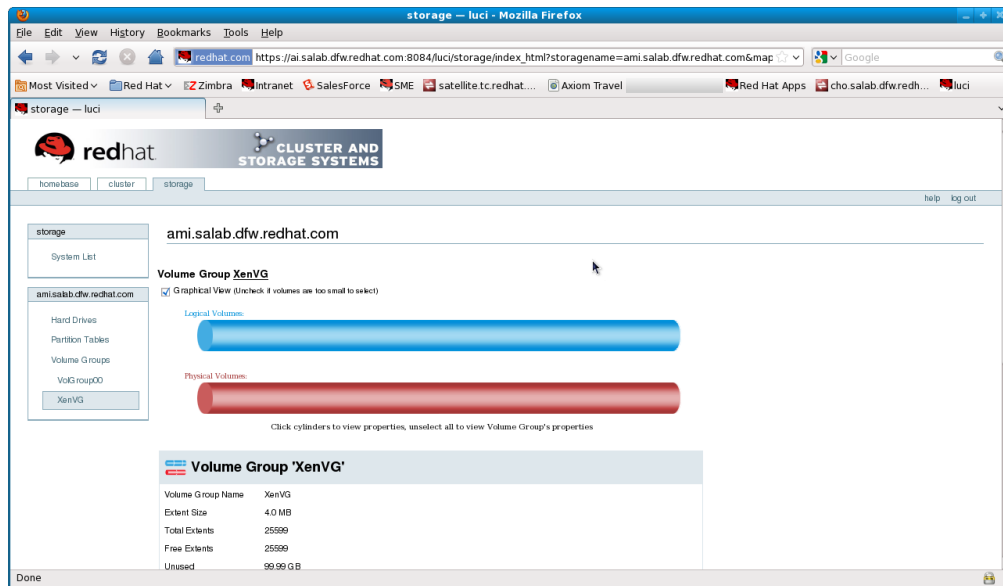












Configure Storage Using Conga

Set up the logical volume which will be used to house the virtual machines

- Scroll down

- New Logical Volume

- Define LV Name

- Define LV Size

- Define Content (GFS2 recommended)

- Define Unique GFS Name (gfs0)

- Define mount point (/var/lib/xen/images)

Configure Storage Using Conga

Set up the logical volume which will be used to house the virtual machines

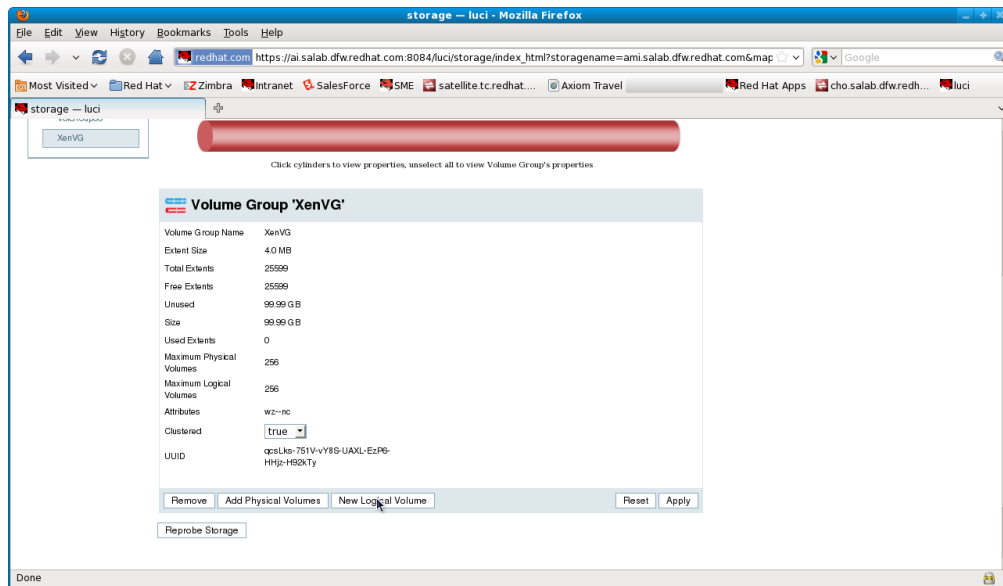
- Set Mount to true

- Set List in /etc/fstab to true

- Set number of journals

- Verify Clustered is set to true

I usually add a couple of spare journals so that if I want to add nodes later there are journals for them. Not nearly as big a deal on GFS2 as on GFS1.



storage — luci - Mozilla Firefox

https://ai.salab.dfw.redhat.com:8084/luci/storage/?pagetype=62&mapper_id=volume_group:XenVG&mapper_id=XenLV

Volume Groups

XenVG

Physical Volumes

Click cylinders to view properties, unselect all to view Volume Group's properties

Unused Space - Creating New Logical Volume

| | | | |
|---------------------|-------------------------|--------------------|----------------------|
| Logical Volume Name | XenLV | Content | GFS2 - Global FS v.2 |
| Volume Group Name | XenVG | Cluster Name | dom0-cluster |
| Size | 99.99 (0.00 - 99.99) GB | Unique GFS Name | gfs0 |
| Clustered | true | Journal Size | 32.0 MB |
| | | Mountpoint | /var/lib/xen/images |
| | | Mount | true |
| | | List in /etc/fstab | true |
| | | Mountable | true |
| | | Locking Protocol | dm |
| | | Number of Journals | 5 (1 - 128) |
| | | Clustered | true |

Reset Create

Done

RED HAT :: CHICAGO :: 2009
SUMMIT

In some rare cases, once you've created the volume group on one host, you might need to restart the `clvmd` service on the other nodes so they register the new volume group. Alternatively, just re-probe storage through the Conga web UI.

```
root@ami:~  
File Edit View Terminal Help  
15289 ? S< 0:00 [iscsi_q_5]  
15290 ? S< 0:00 [scsi_wq_5]  
19327 ? Rs 0:00 sshd: root@pts/1  
19329 pts/1 Ss 0:00 -bash  
22497 ? Ssl 0:00 /sbin/ccsd  
22507 ? Sll 0:00 aisexec  
22544 ? Ss 0:00 /sbin/groupd  
22551 ? Ss 0:00 /sbin/fenced  
22557 ? Ss 0:00 /sbin/dlm_controld  
22563 ? Ss 0:00 /sbin/gfs_controld  
22583 ? Ssl 0:00 clvmd -T20  
22584 ? S< 0:00 [dlm_astd]  
22585 ? S< 0:00 [dlm_sndd]  
22586 ? S< 0:00 [dlm_rcv]  
22587 ? S< 0:00 [dlm_send]  
22588 ? S< 0:00 [dlm_recoverd]  
22641 ? S<Ls 0:00 clurgmgrd  
22643 ? S<L 0:00 clurgmgrd  
22645 ? S< 0:00 [dlm_recoverd]  
23175 ? S<S 0:00 /usr/libexec/ricci/ricci-worker -f /var/lib/ricci/que  
23176 ? S 0:00 ricci_modstorage  
23266 ? R 0:00 /sbin/mkfs.gfs2 -J 32 -j 5 -O -p lock_dlm -t dom0-clu  
23269 pts/1 R+ 0:00 ps ax  
[root@ami ~]#
```

After submitting the job, you can watch the mkfs.gfs2 process running on the target machine.

Configure Storage Using Conga

Add the mount point to the other hosts in the cluster

Storage

System List

System name

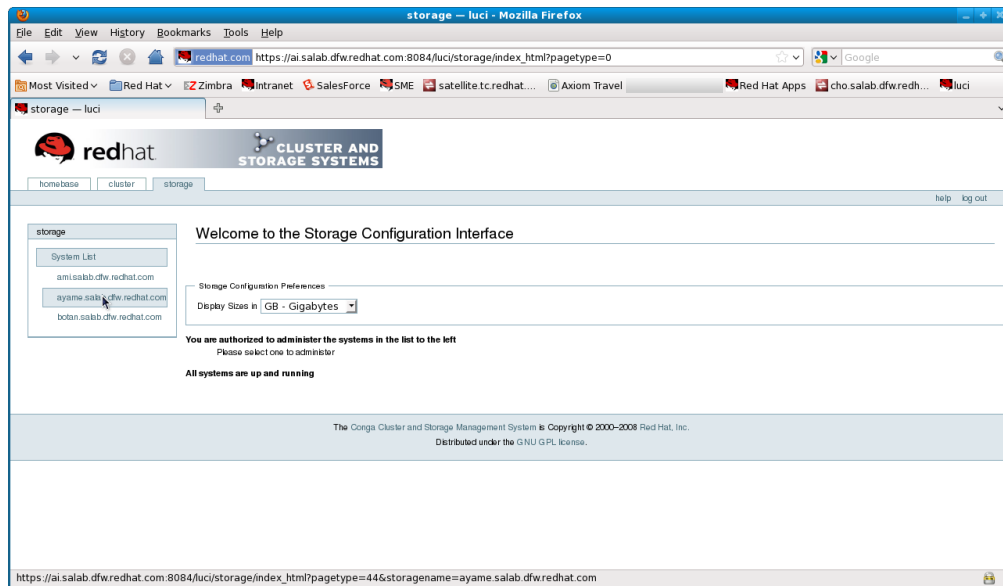
Expand VG

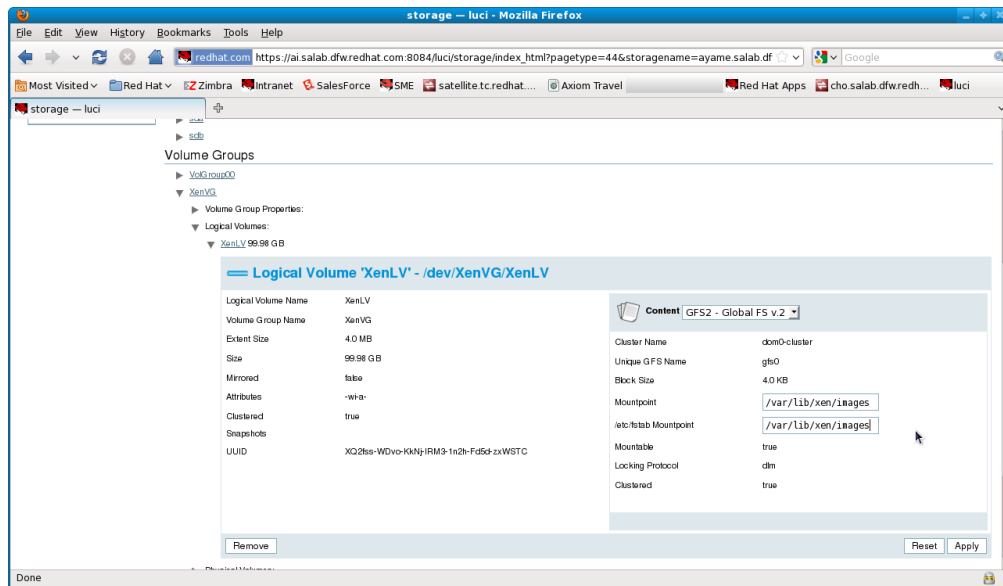
Expand LV

Fill in Mountpoint

Fill in /etc/fstab Mountpoint

In some cases, you might need to reprobe storage so that the host “sees” the new LV





Configure Storage (alternative method)

Copy and paste the last line of /etc/fstab to the other nodes and then mount -a

```
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]# mount  
/dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/sdal on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /var/lib/xenstored type tmpfs (rw)  
none on /sys/kernel/config type configfs (rw)  
[root@botan ~]# echo /dev/XenVG/XenLV /var/lib/xen/images gfs2 defaults 0 0 >> /  
etc/fstab  
[root@botan ~]# mount -a  
[root@botan ~]# mount  
/dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/sdal on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /var/lib/xenstored type tmpfs (rw)  
none on /sys/kernel/config type configfs (rw)  
/dev/mapper/XenVG-XenLV on /var/lib/xen/images type gfs2 (rw,hostdata=jid=2:id=1  
31073:first=0)  
[root@botan ~]#
```


Verify Storage Availability

Something simple like a `df -h` on all nodes is fine.

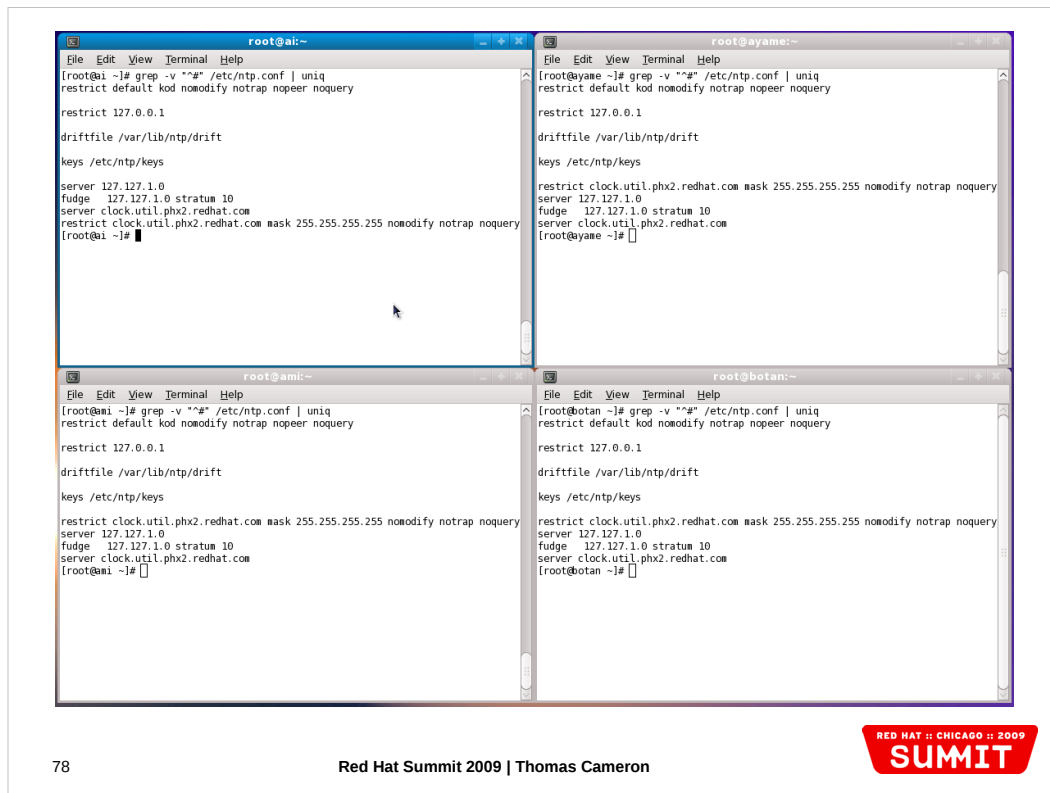
The image displays four terminal windows, each showing the output of the `df -h` command for a different user. The output is consistent across all users, indicating a shared filesystem structure.

| Filesystem | Size | Used | Avail | Use% | Mounted on |
|---------------------------------|------|------|-------|------|---------------------|
| /dev/mapper/VolGroup00-LogVol00 | 901G | 2.9G | 851G | 1% | / |
| /dev/sdal | 99M | 48M | 46M | 52% | /boot |
| tmpfs | 3.7G | 0 | 3.7G | 0% | /dev/shm |
| none | 3.7G | 104K | 3.7G | 1% | /var/lib/xenstored |
| /dev/mapper/XenVG-XenLV | 100G | 166M | 100G | 1% | /var/lib/xen/images |

Prepare dom0 machines for live migration

Make sure that time is synced across all of the dom0 machines.

You **must** use something like NTP. Even if you just set up the management server as an NTP server, the dom0 times **must** be synced for live migration to work right!



In the lab, all of the servers are syncing against our internal time source. All of the servers time is synced.

Install Virtual Machines

Use virt-manager or virt-install

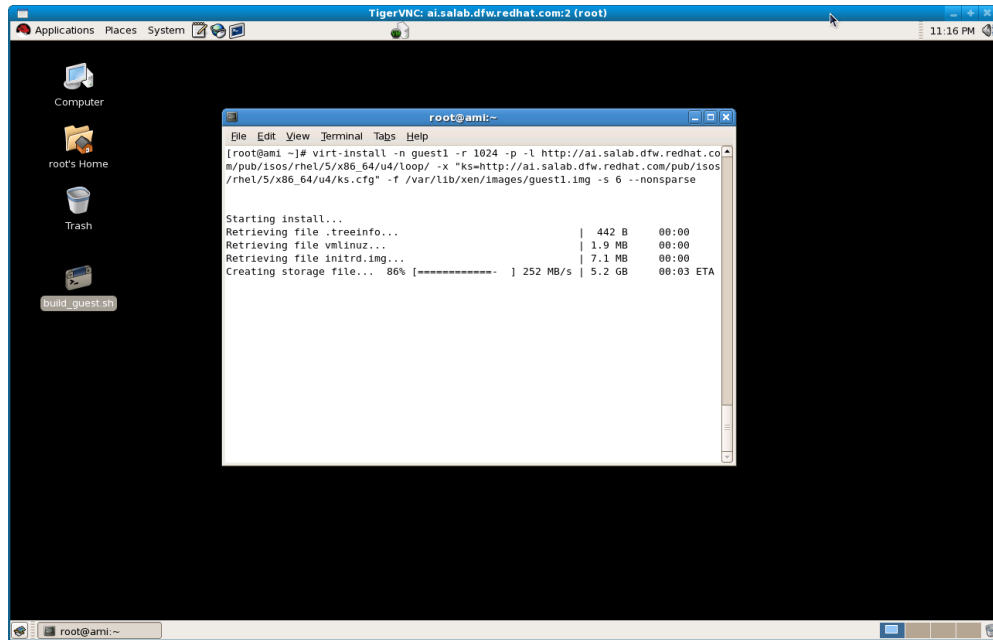
The following is a single command, separated here for easier viewing:

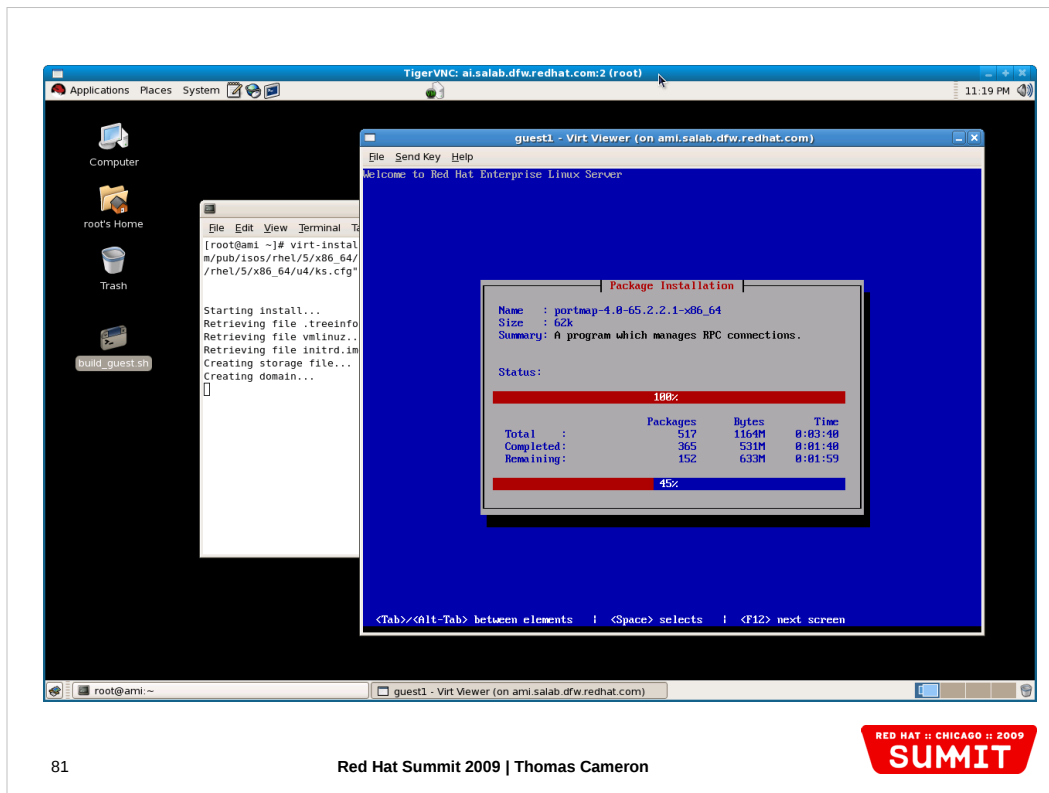
```
virt-install -n guest1 -r 1024 -p \  
-l http://ai.salab.dfw.redhat.com/pub/isos/rhel/5/x86_64/u4/loop/ \  
-x "ks=http://ai.salab.dfw.redhat.com/pub/isos/rhel/5/x86_64/u4/ks.cfg" \  
-f /var/lib/xen/images/guest1.img -s 6 --nonsparse  
~~~~~
```

In my lab, I created two domU guests on each dom0 host.

--non-sparse is very important for performance!

For mass provisioning via virt-install we also added an option to not wait in the foreground for the install to complete '--wait'





Since I am using a single gigabit NIC for each dom0, I only installed one guest at a time per dom0. The iSCSI target, ai, was pretty loaded with three concurrent installations going but it ran pretty well. For a production system I would definitely do bonded ethernet on the dom0 and the iSCSI target.

Prepare all dom0 machines for live migration

In the `/etc/xend-config.sxp` file set the following:

```
(xend-relocation-server yes)
(xend-relocation-port 8002)
(xend-relocation-address "")
(xend-relocation-hosts-allow "")
```

Restart the xend service

Test live migration

```
virsh migrate guest1 \ xen
+ssh://root@ayame.salab.dfw.redhat.com
```

Note that these are two single quotes (') not a single double-quote

For testing migration, since virsh uses ssh, you might want to create ssh keys. See the man page for `ssh-keygen(1)`

Note that this is only an example and done in a secure lab. Consider security implications of opening up your server this wide!


```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# grep xend-relocation /etc/xen/xend-config.sxp | grep -v "^#"  
(xend-relocation-server yes)  
(xend-relocation-port 8002)  
(xend-relocation-address '')  
(xend-relocation-hosts-allow '')  
[root@ami ~]# service xend restart  
restart xend: [ OK ]  
[root@ami ~]# virsh migrate guest1 xen+ssh://root@ayame.salab.dfw.redhat.com  
root@ayame.salab.dfw.redhat.com's password:  
[root@ami ~]# virsh --connect xen+ssh://root@ayame.salab.dfw.redhat.com  
root@ayame.salab.dfw.redhat.com's password:  
Welcome to virsh, the virtualization interactive terminal.  
  
Type: 'help' for help with commands  
      'quit' to quit  
  
virsh # list  
Id Name State  
-----  
0 Domain-0 running  
4 guest1 idle  
virsh #
```

Prepare dom0 machines for live migration

Stop all the guests

```
chkconfig xendomains off
```

```
service xendomains stop
```

We don't want the xendomains service to try to start a guest at the same time that the clustering services are trying to start them, so we'll turn all start/stop abilities over the clustering



```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# virsh list  
Id Name State  
-----  
0 Domain-0 running  
  
[root@ami ~]# chkconfig xendomains off  
[root@ami ~]# service xendomains stop  
Shutting down Xen domains:[done] [ OK ]  
[root@ami ~]#
```

Prepare dom0 machines for live migration

Normally you would move the config files etc /etc/xen to shared storage (/var/lib/xen/images on each dom0), but there is a bug in the beta version of RHEL 5.4, BZ 519786.

Either:

- use the vm.sh from the BZ

- copy all of the domU config files to /etc/xen on all of the nodes

- symlink the configs from /etc/xen/[guest] to /var/lib/xen/images/[guest] on all the nodes.

For this presentation, I used the vm.sh from the BZ and moved all the configs to shared storage

Since the cluster service needs to have access to the config files, they need to be on shared storage. Since /var/lib/xen/images is already shared, just move them there.



```
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# ls /var/lib/xen/images/  
guest1.img guest2.img guest3.img guest4.img guest5.img guest6.img  
[root@ami ~]# mv /etc/xen/guest* /var/lib/xen/images/  
[root@ami ~]# ls /var/lib/xen/images/  
guest1    guest2    guest3    guest4    guest5    guest6  
guest1.img guest2.img guest3.img guest4.img guest5.img guest6.img  
[root@ami ~]#
```

This shows the results of having moved the config files on all three dom0 machines at the same time.

Configure Failover Domains

Cluster tab

Cluster Name

Failover Domains

Add Failover Domain

FD Name

Check "Prioritized"

Check "Restrict failover to this domain's members"

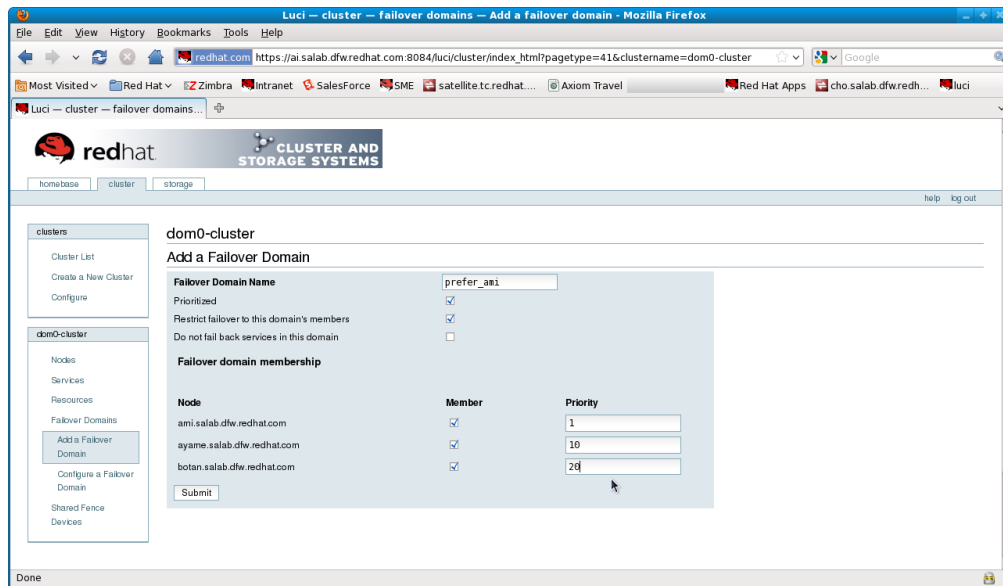
Do not check "Do not fail back services..."

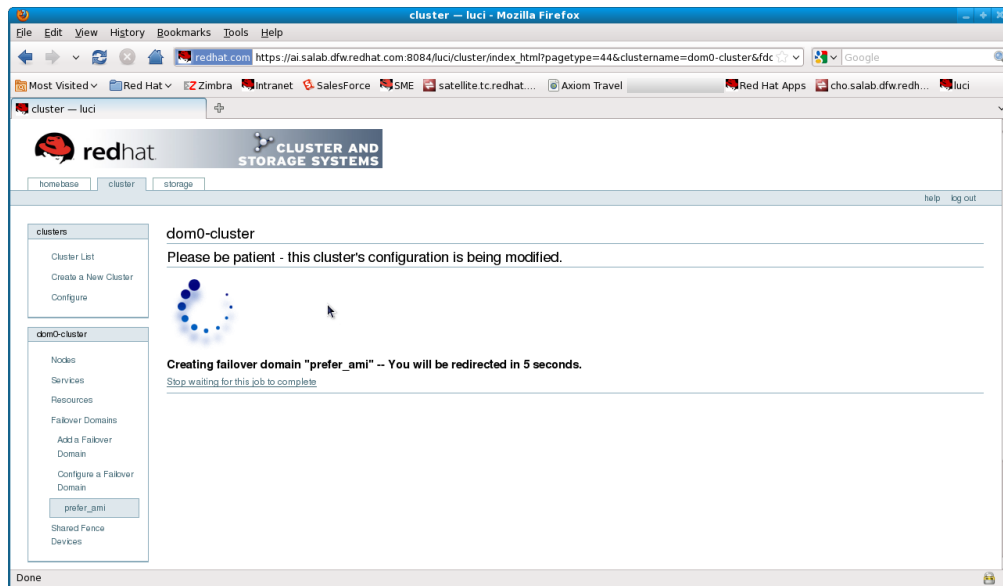
Assign members and priorities

Submit

Prioritized sets services assigned to the FD to run on the lowest priority first

Set up FD for each server in the dom0 cluster.





Configure Failover Domains

Create a domain for each dom0.

prefer_ami

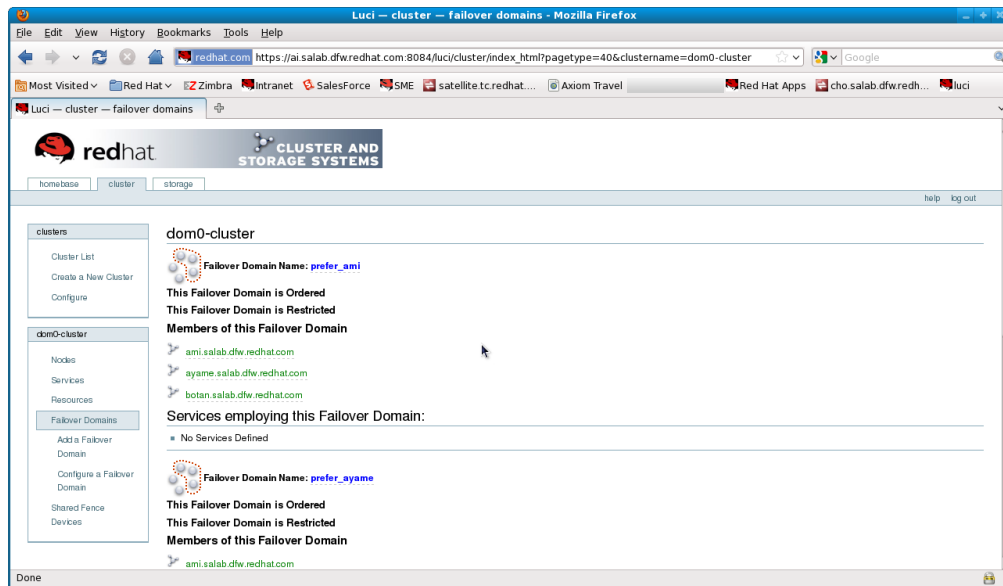
ami 1, ayame 10, botan 20

prefer_ayame

ayame 1, botan 10, ami 20

prefer_botan

botan 1, ami 10, ayame 20



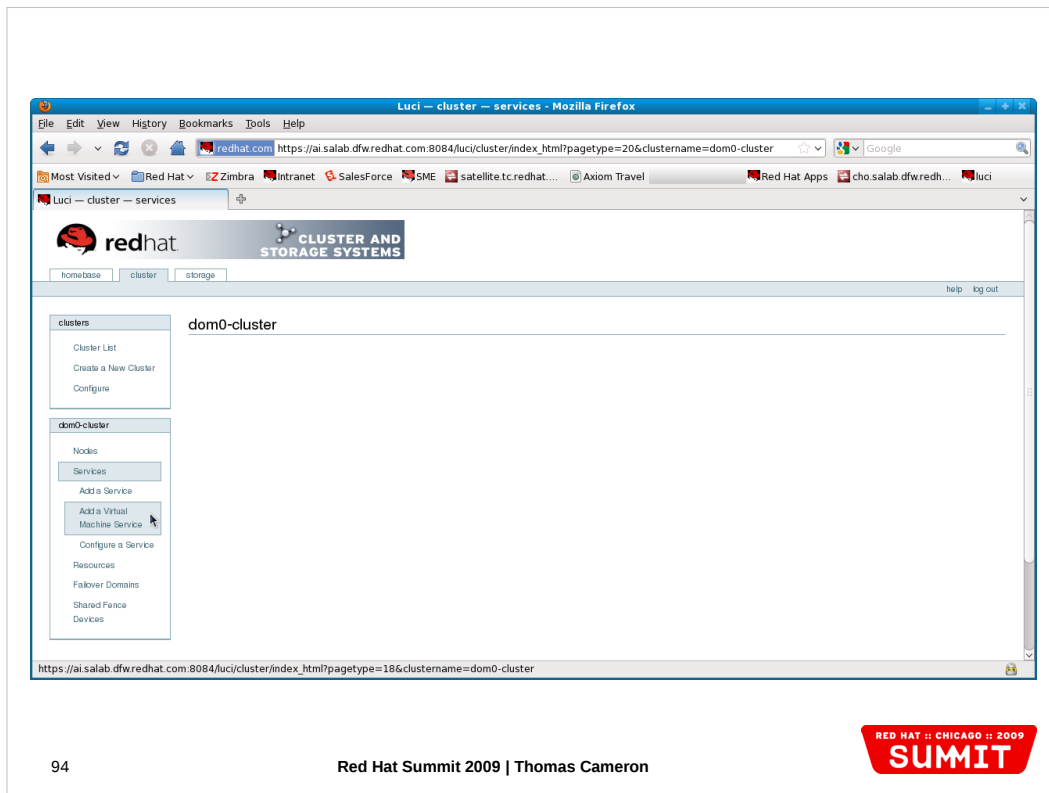
Define Virtual Machine Services

Cluster Tab

Cluster Name

Services

Add Virtual Machine Service



Note BZ 519252 - Conga does not show the option to add a virtual machine service. Workaround is to edit `/var/lib/luci/Extensions/ricci_communicator.py` and change these lines:

```
self.__dom0 = hello.firstChild.getAttribute('xen_host') == 'true'
```

to:

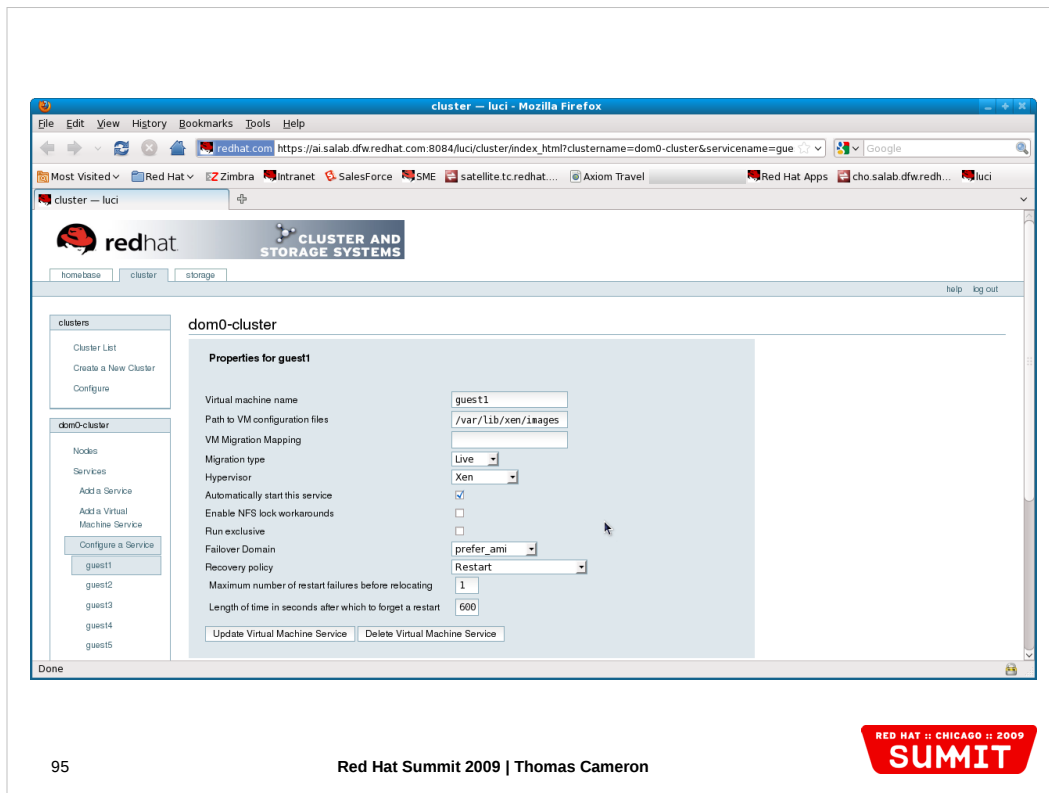
```
self.__dom0 = True
```

```
self.__dom0 = resp.firstChild.getAttribute('xen_host') == 'true'
```

to:

```
self.__dom0 = True
```

Also, FD will not show up correctly in luci-0.12.2-6.el5. Update to luci-0.12.2-6.el5 from RHEL 5.4 to fix it.



VM name must be the same as the config file name

VM Migration Mapping is for multi-homed machines, if you want to specify a management network interface

HV should be set to Xen

No NFS so skip it

Run exclusive allows no other services on this cluster, skip it

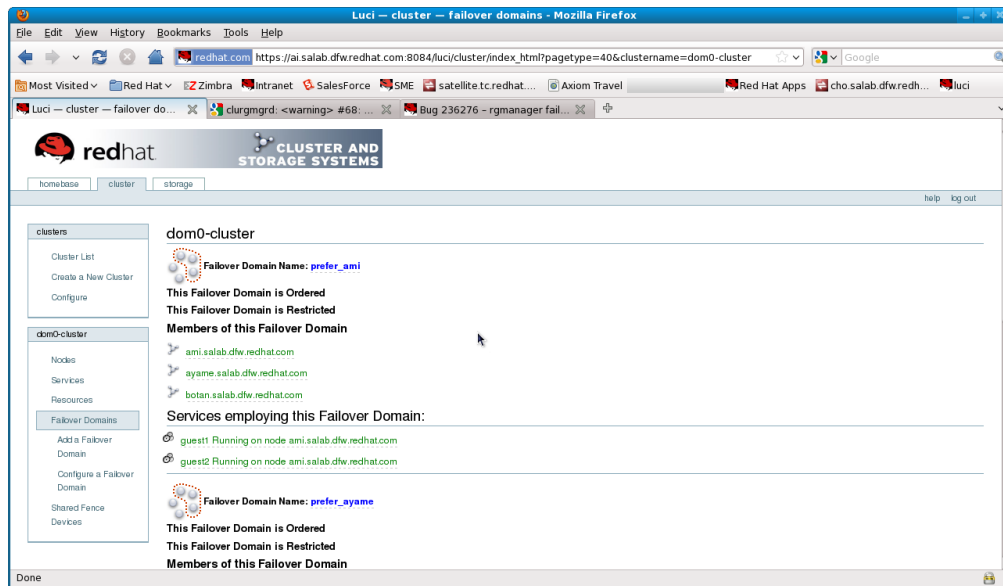
Example recovery policy sets restart unless it fails more than once in 10 minutes, then relo

If you have problems getting the VMs to come up in the cluster, manually start them and then activate them in the Conga interface.

Define Virtual Machine Services

Repeat for each VM Service

In this example, two VM Services are assigned to each failover domain - this equates to two machines starting on ai, two on ayame and two on botan



Test virtual machine management services

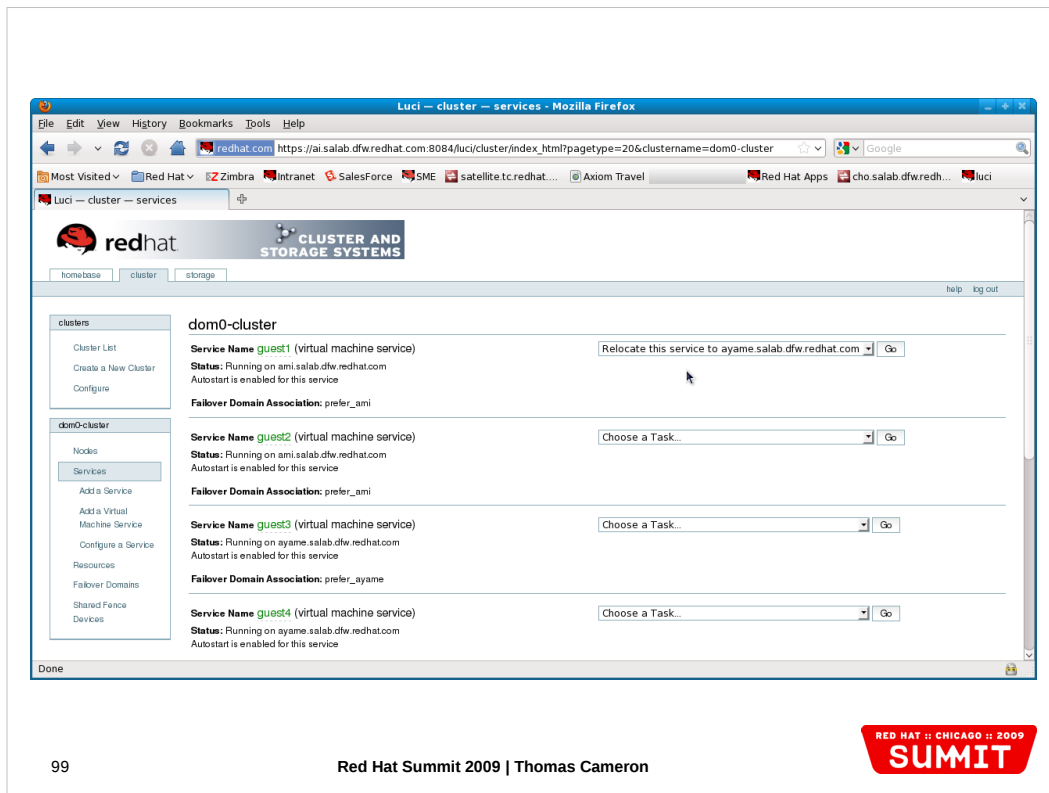
Cluster Tab

Cluster Name

Services

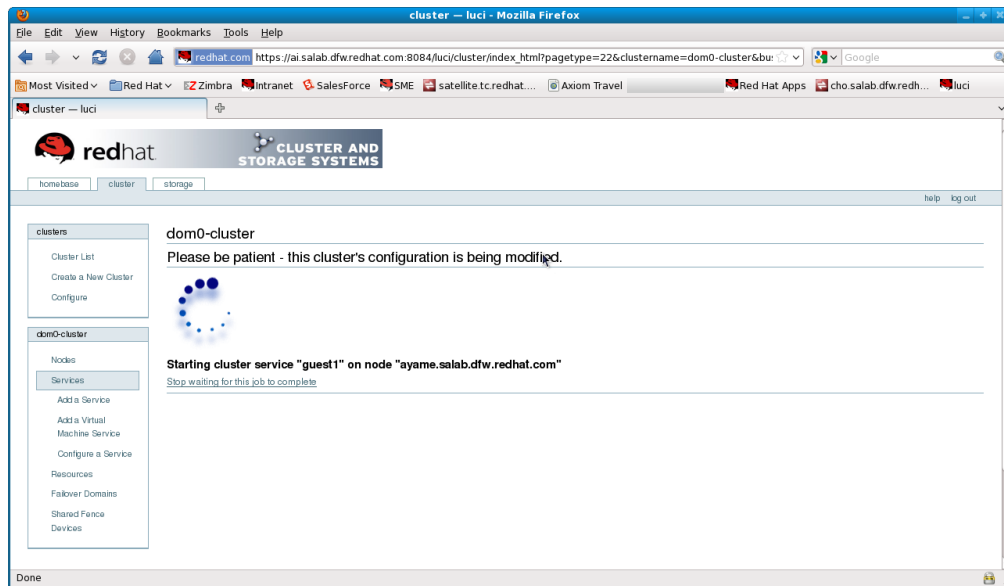
Choose an action for one of the machines from the drop down menu

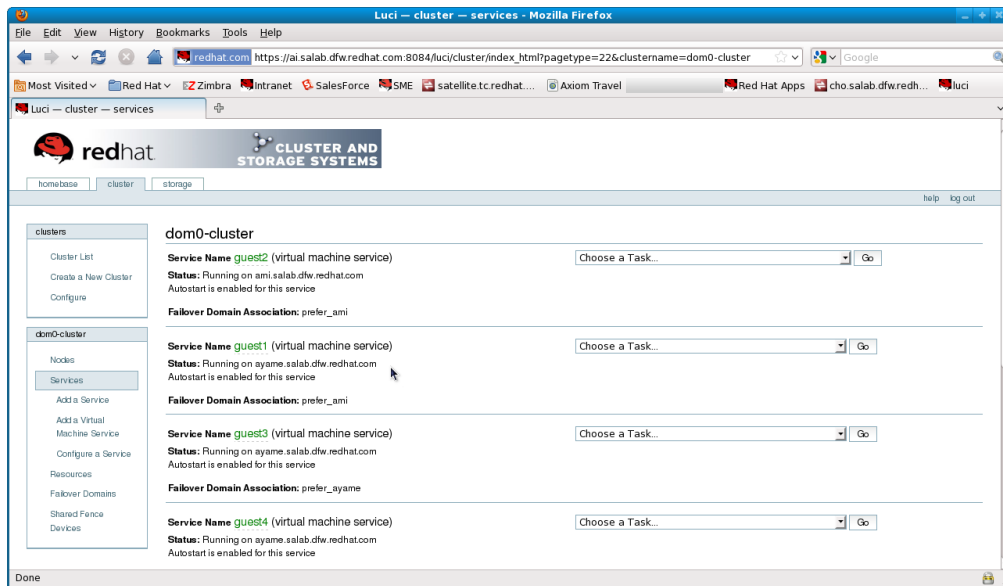
Go and confirm



You can also do this from the command line:

```
[root@ayame ~]# clusvcadm -M vm:guest1 -m ayame
'ayame' not in membership list
Closest match: 'ayame.salab.dfw.redhat.com'
Trying to migrate vm:guest1 to
ayame.salab.dfw.redhat.com...Success
[root@ayame ~]#
```





```
root@ayame:~  
File Edit View Terminal Help  
[root@ayame ~]# virsh list  
Id Name State  
-----  
0 Domain-0 running  
5 guest4 idle  
6 guest3 idle  
8 guest1 idle  
  
[root@ayame ~]# clustat  
Cluster Status for dom0-cluster @ Thu Aug 27 17:10:35 2009  
Member Status: Quorate  
  
Member Name ID Status  
-----  
am1.salab.dfw.redhat.com 1 Online, rgmanager  
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com 3 Online, rgmanager  
  
Service Name Owner (Last) State  
-----  
vm:guest1 ayame.salab.dfw.redhat.com started  
vm:guest2 am1.salab.dfw.redhat.com started  
vm:guest3 ayame.salab.dfw.redhat.com started  
vm:guest4 ayame.salab.dfw.redhat.com started  
vm:guest5 botan.salab.dfw.redhat.com started  
vm:guest6 botan.salab.dfw.redhat.com started  
[root@ayame ~]#
```

Test virtual machine failover services

Log into a domU machine and crash it

```
echo c > /proc/sysrq-trigger
```

Watch clustat on one of the dom0 machines

In the following slides, the domU system is crashed and the cluster restarts it on ami.


```
root@sayoko:~# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 10:56:25 2009
root@sayoko:~# echo c > /proc/sysrq-trigger

root@ayame:~# clustat
Every 2.0s: clustat
Thu Aug 27 17:26:58 2009
Cluster Status for dom0-cluster @ Thu Aug 27 17:26:58 2009
Member Status: Quorate
Member Name ID Status
-----
ami.salab.dfw.redhat.com 1 Online, rgmanager
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager
botan.salab.dfw.redhat.com 3 Online, rgmanager

Service Name Owner (Last) State
-----
va:guest1 (ayame.salab.dfw.redhat.com) recoverable
va:guest2 ami.salab.dfw.redhat.com started
va:guest3 ayame.salab.dfw.redhat.com started
va:guest4 ayame.salab.dfw.redhat.com started
va:guest5 botan.salab.dfw.redhat.com started
va:guest6 botan.salab.dfw.redhat.com started

root@ami:~# virsh console guest1
Connected to domain guest1
Escape character is ^]

Red Hat Enterprise Linux Server release 5.4 (Tikanga)
Kernel 2.6.18-164.el5xen on an x86_64

sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump
Kexec: Warning: crash image not loaded
Kernel panic - not syncing: SysRq-triggered panic!

root@ami:~#

root@botan:~#
```



```
root@sayoko:~  
File Edit View Terminal Help  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 10:56:25 2009  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
[  
root@ayame:~  
File Edit View Terminal Help  
Every 2.0s: clustat Thu Aug 27 17:27:04 2009  
Cluster Status for dom0-cluster @ Thu Aug 27 17:27:04 2009  
Member Status: Quorate  
Member Name ID Status  
-----  
ami.salab.dfw.redhat.com 1 Online, rgmanager  
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com 3 Online, rgmanager  
Service Name Owner (Last) State  
-----  
vm:guest1 ami.salab.dfw.redhat.com started  
vm:guest2 ami.salab.dfw.redhat.com started  
vm:guest3 ayame.salab.dfw.redhat.com started  
vm:guest4 ayame.salab.dfw.redhat.com started  
vm:guest5 botan.salab.dfw.redhat.com started  
vm:guest6 botan.salab.dfw.redhat.com started  
root@ami:~  
File Edit View Terminal Help  
[root@ami ~]# virsh console guest1  
Connected to domain guest1  
Escape character is ^]  
Red Hat Enterprise Linux Server release 5.4 (Tikanga)  
Kernel 2.6.18-164.el5xen on an x86_64  
sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump  
Kexec: Warning: crash image not loaded  
Kernel panic - not syncing: SysRq-triggered panic!  
[root@ami ~]#  
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]#
```

Test virtual machine failover services

The same guest is crashed a second time and the cluster migrates it to ayami based on the failover domain rules we set up earlier.

```
root@sayoko:~# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 10:56:25 2009
[root@sayoko ~]# echo c > /proc/sysrq-trigger
Read from remote host sayoko.salab.dfw.redhat.com: Connection reset by peer
Connection to sayoko.salab.dfw.redhat.com closed.
[root@ai ~]#
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 17:23:27 2009 from ai.salab.dfw.redhat.com
[root@sayoko ~]# echo c > /proc/sysrq-trigger

root@ayame:~#
Every 2.0s: clustat Thu Aug 27 17:27:54 2009
Cluster Status for dom0-cluster @ Thu Aug 27 17:27:54 2009
Member Status: Quorate

Member Name ID Status
-----
ani.salab.dfw.redhat.com 1 Online, rgmanager
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager
botan.salab.dfw.redhat.com 3 Online, rgmanager

Service Name Owner (Last) State
-----
vm:quest1 ani.salab.dfw.redhat.com started
vm:quest2 ani.salab.dfw.redhat.com started
vm:quest3 ayame.salab.dfw.redhat.com started
vm:quest4 ayame.salab.dfw.redhat.com started
vm:quest5 botan.salab.dfw.redhat.com started
vm:quest6 botan.salab.dfw.redhat.com started

root@ami:~#
device-mapper: uevent: version 1.0.3
device-mapper: ioctl: 4.11.5-ioctl (2007-12-12) initialised: dm-devel@redhat.com
device-mapper: dm-raid45: initialized v0.25941
EXT3-fs: INFO: recovery required on readonly filesystem.
EXT3-fs: write access will be enabled during recovery.
kjournald starting. Commit interval 5 seconds
EXT3-fs: recovery complete.
EXT3-fs: mounted filesystem with ordered data mode.
SELinux: Disabled at runtime.
type=1404 audit(1251412033.542:2): selinux=0 auid=4294967295 ses=4294967295
Bridge firewalling registered
ip_tables: (C) 2000-2006 Netfilter Core Team
Netfilter messages via NETLINK v0.30.
ip_conntrack version 2.4 (8192 buckets, 65536 max) - 304 bytes per conntrack

Red Hat Enterprise Linux Server release 5.4 (Tikanga)
Kernel 2.6.18-164.el5xen on an x86_64

sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump
kexec: Warning: crash image not loaded
Kernel panic - not syncing: SysRq-triggered panic!
```



```
root@sayoko:~  
File Edit View Terminal Help  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 10:56:25 2009  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
Read from remote host sayoko.salab.dfw.redhat.com: Connection reset by peer  
Connection to sayoko.salab.dfw.redhat.com closed.  
[root@ai ~]#  
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com  
Last login: Thu Aug 27 17:23:27 2009 from ai.salab.dfw.redhat.com  
[root@sayoko ~]# echo c > /proc/sysrq-trigger  
[root@sayoko ~]#  
root@ayame:~  
File Edit View Terminal Help  
Every 2.0s: clustat Thu Aug 27 17:28:08 2009  
Cluster Status for dom0-cluster @ Thu Aug 27 17:28:08 2009  
Member Status: Quorate  
Member Name ID Status  
-----  
ami.salab.dfw.redhat.com 1 Online, rgmanager  
ayame.salab.dfw.redhat.com 2 Online, Local, rgmanager  
botan.salab.dfw.redhat.com 3 Online, rgmanager  
Service Name Owner (Last) State  
-----  
vm:quest1 (ami.salab.dfw.redhat.com) recoverable  
vm:quest2 ami.salab.dfw.redhat.com started  
vm:quest3 ayame.salab.dfw.redhat.com started  
vm:quest4 ayame.salab.dfw.redhat.com started  
vm:quest5 botan.salab.dfw.redhat.com started  
vm:quest6 botan.salab.dfw.redhat.com started  
root@ami:~  
File Edit View Terminal Help  
device-mapper: uevent: version 1.0.3  
device-mapper: ioctl: 4.11.5-ioctl (2007-12-12) initialised: dm-devel@redhat.com  
device-mapper: dm-raid45: initialized v0.25941  
EXT3-fs: INFO: recovery required on readonly filesystem.  
EXT3-fs: write access will be enabled during recovery.  
kjournald starting. Commit interval 5 seconds  
EXT3-fs: recovery complete.  
EXT3-fs: mounted filesystem with ordered data mode.  
SELinux: Disabled at runtime.  
type=1404 audit(1251412033.542:2): selinux=0 audit=4294967295 ses=4294967295  
Bridge firewalling registered  
ip_tables: (C) 2000-2006 Netfilter Core Team  
Netfilter messages via NETLINK v0.30.  
ip_conntrack version 2.4 (8192 buckets, 65536 max) - 304 bytes per conntrack  
Red Hat Enterprise Linux Server release 5.4 (Tikanga)  
Kernel 2.6.18-164.el5xen on an x86_64  
sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump  
kexec: Warning: crash image not loaded  
Kernel panic - not syncing: SysRq-triggered panic!  
[root@ami ~]#  
root@botan:~  
File Edit View Terminal Help  
[root@botan ~]#
```



```
root@sayoko:~# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 10:56:25 2009
[root@sayoko ~]# echo c > /proc/sysrq-trigger
Read from remote host sayoko.salab.dfw.redhat.com: Connection reset by peer
Connection to sayoko.salab.dfw.redhat.com closed.
[root@ai ~]#
[root@ai ~]# ssh sayoko.salab.dfw.redhat.com
Last login: Thu Aug 27 17:23:27 2009 from ai.salab.dfw.redhat.com
[root@sayoko ~]# echo c > /proc/sysrq-trigger

root@ayame:~#
Every 2.0s: clustat                               Thu Aug 27 17:28:12 2009
Cluster Status for dom0-cluster @ Thu Aug 27 17:28:13 2009
Member Status: Quorate

Member Name                                ID    Status
-----
ami.salab.dfw.redhat.com                   1     Online, rgmanager
ayame.salab.dfw.redhat.com                 2     Online, Local, rgmanager
botan.salab.dfw.redhat.com                 3     OnLine, rgmanager

Service Name                                Owner (Last)                                State
-----
va:guest1                                  none                                         recovering
va:guest2                                  ami.salab.dfw.redhat.com                  started
va:guest3                                  ayame.salab.dfw.redhat.com                started
va:guest4                                  ayame.salab.dfw.redhat.com                started
va:guest5                                  botan.salab.dfw.redhat.com                started
va:guest6                                  botan.salab.dfw.redhat.com                started

root@ami:~#
device-mapper: uevent: version 1.0.3
device-mapper: ioctl: 4.11.5-ioctl (2007-12-12) initialised: dm-devel@redhat.com
device-mapper: dm-raid45: initialized v0.25941
EXT3-fs: INFO: recovery required on readonly filesystem.
EXT3-fs: write access will be enabled during recovery.
kjournald starting. Commit interval 5 seconds
EXT3-fs: recovery complete.
EXT3-fs: mounted filesystem with ordered data mode.
SELinux: Disabled at runtime.
type=1404 audit(1251412033.542:2): selinux=0 audit=4294967295 ses=4294967295
Bridge firewalling registered
ip_tables: (C) 2000-2006 Netfilter Core Team
Netfilter messages via NETLINK v0.30.
ip_conntrack version 2.4 (8192 buckets, 65536 max) - 304 bytes per conntrack
Red Hat Enterprise Linux Server release 5.4 (Tikanga)
Kernel 2.6.18-164.el5xen on an x86_64

sayoko.salab.dfw.redhat.com login: SysRq : Trigger a crashdump
Kexec: Warning: crash image not loaded
Kernel panic - not syncing: SysRq-triggered panic!

root@botan:~#
```



Set up the clustered app on the domU cluster

Define a new cluster - domU-cluster

- Cluster tab

- Create new cluster

- Cluster name

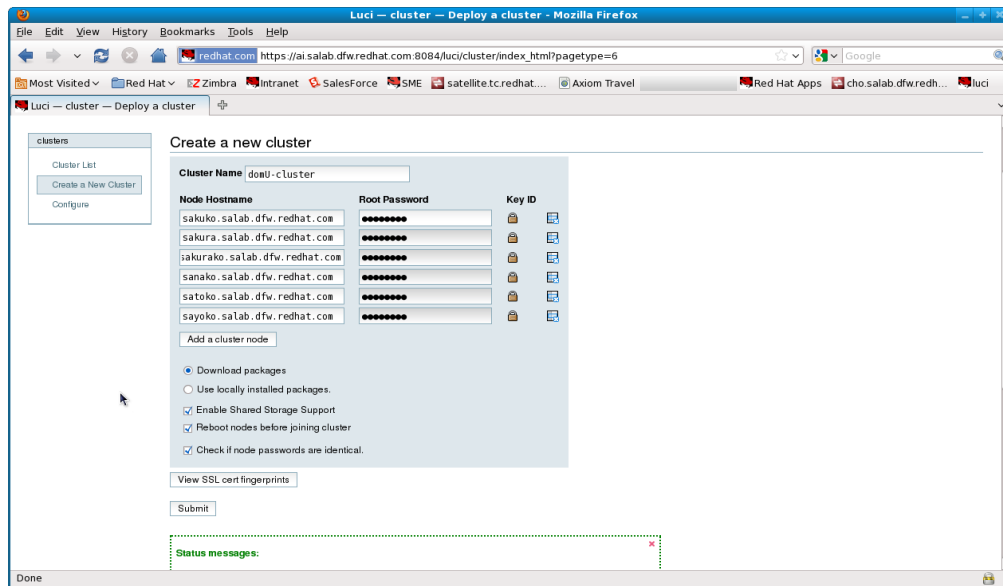
- Node Hostnames/passwords

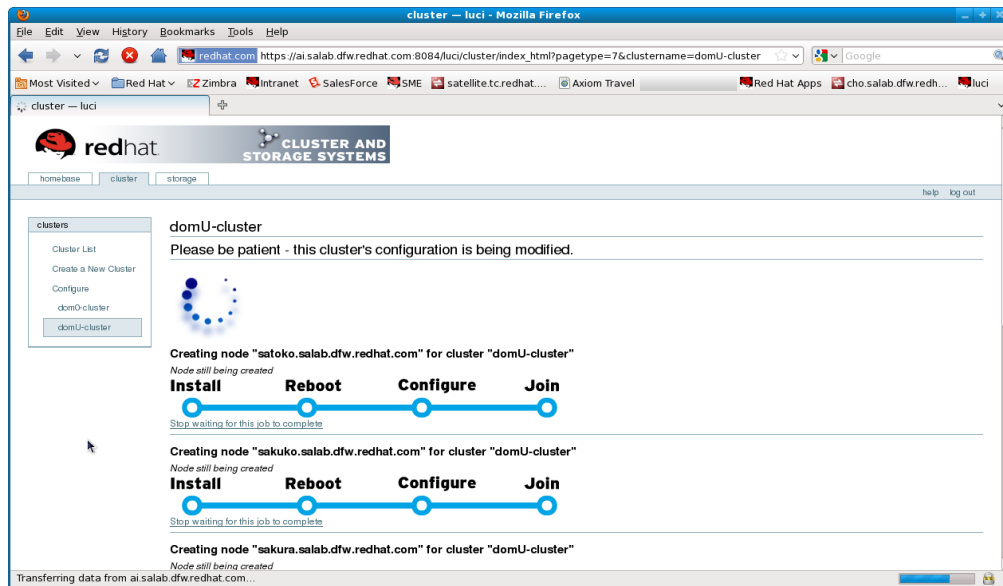
- Enable share storage

- Reboot nodes

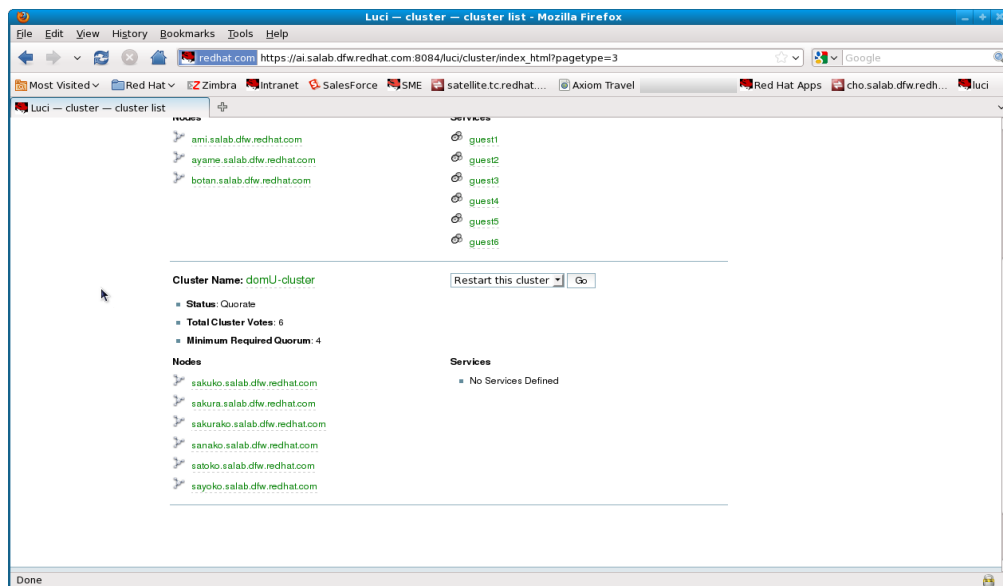
- Check if passwords identical

- View SSL fingerprints





```
root@ayame:~  
File Edit View Terminal Help  
vm:guest2          ami.salab.dfw.redhat.com  started  
vm:guest3          (none)                   recoverable  
vm:guest4          (none)                   recoverable  
vm:guest5          botan.salab.dfw.redhat.com started  
vm:guest6          (none)                   recoverable  
[root@ayame ~]# clustat  
Cluster Status for dom0-cluster @ Fri Aug 28 13:56:35 2009  
Member Status: Quorate  
  
Member Name          ID  Status  
-----  
ami.salab.dfw.redhat.com      1  Online, rgmanager  
ayame.salab.dfw.redhat.com    2  Online, Local, rgmanager  
botan.salab.dfw.redhat.com    3  Online, rgmanager  
  
Service Name          Owner (Last)          State  
-----  
vm:guest1             ami.salab.dfw.redhat.com started  
vm:guest2             (none)                recoverable  
vm:guest3             (none)                recoverable  
vm:guest4             (none)                recoverable  
vm:guest5             botan.salab.dfw.redhat.com started  
vm:guest6             (none)                recoverable  
[root@ayame ~]#
```



Set up the clustered app on the domU cluster

Copy the fence_xvm.key to the domU cluster members

Either scp or use the web UI

Set up the clustered app on the domU cluster

Set up a shared fence device

- Cluster Tab

- Cluster Name

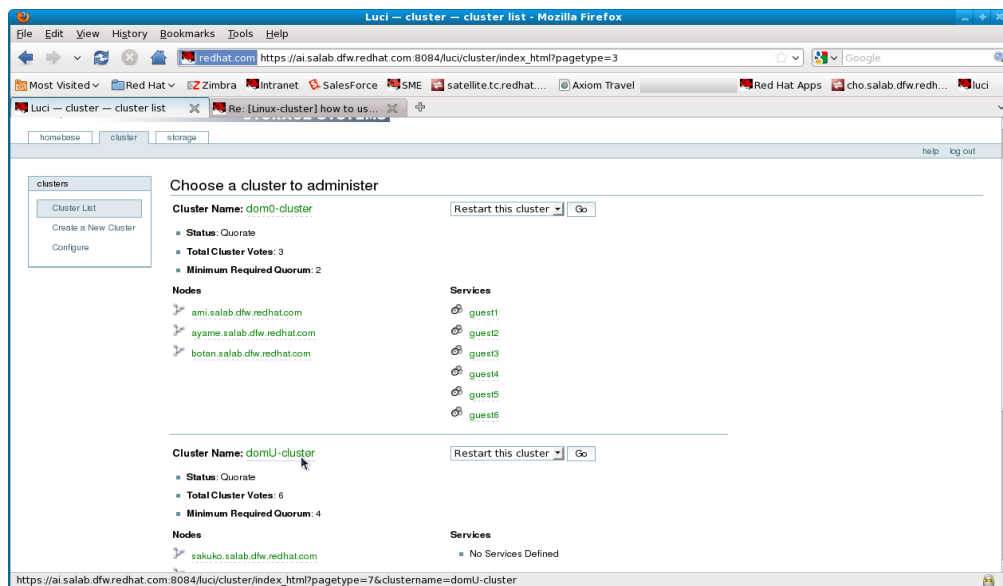
- Shared Fence Devices

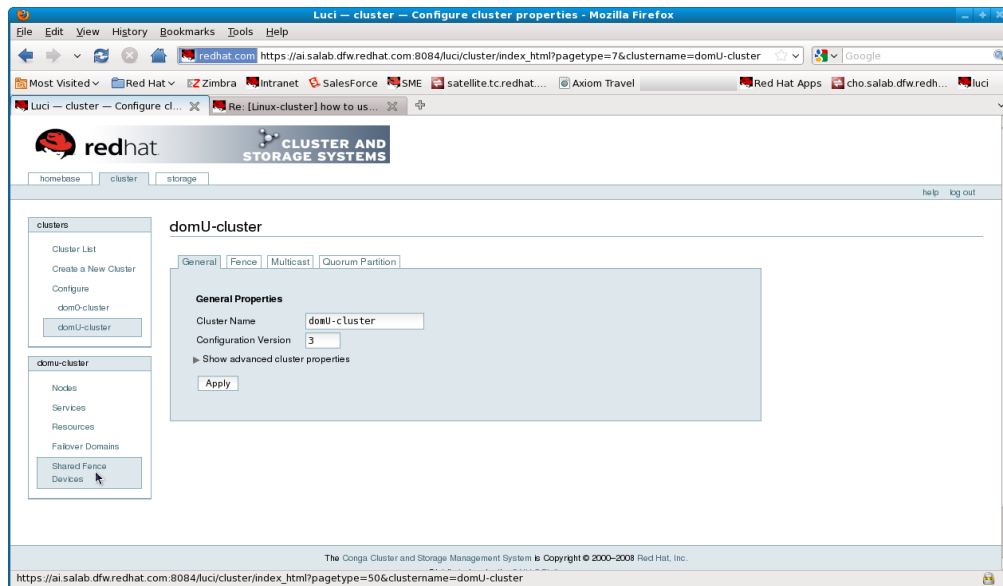
- Add a Fence Device

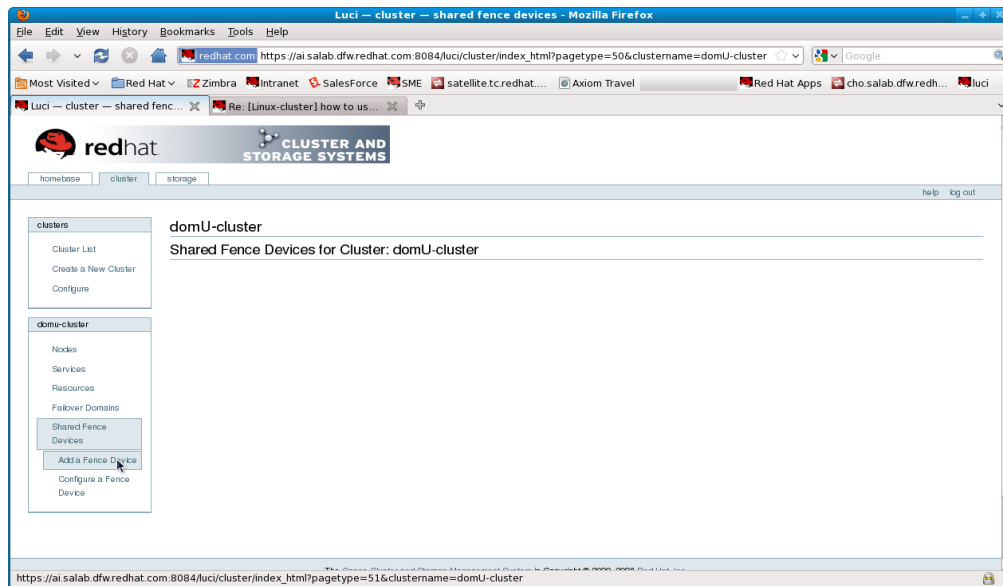
 - Type

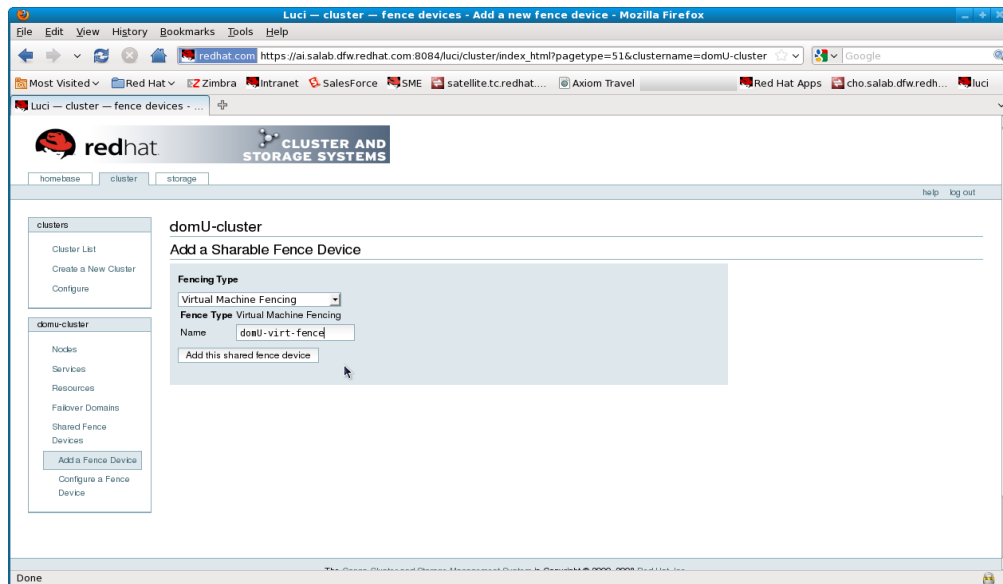
 - Name

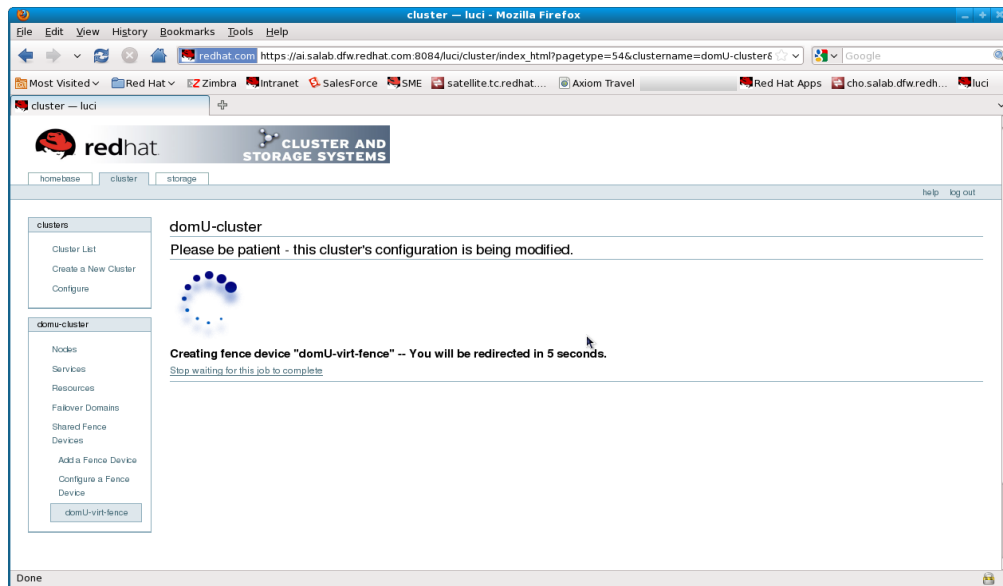
- Submit and confirm











Set up the clustered app on the domU cluster

Assign Fence Devices to Each domU

- Cluster Node

- Main Fencing Method

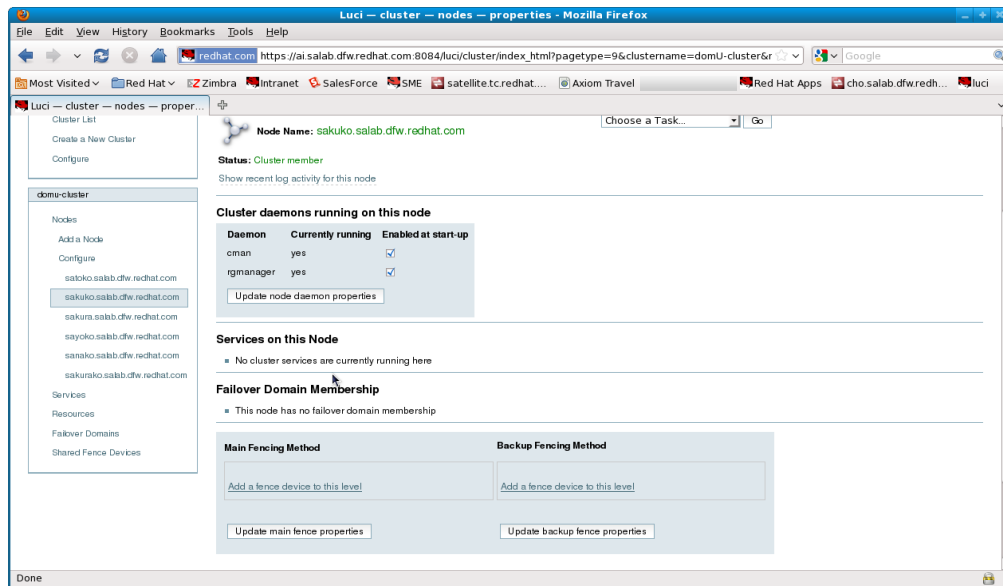
- Add Fence Device To This Level

 - Pre-defined device

 - Domain name

 - Update and confirm

- Repeat for each domU machine



Luci - cluster - nodes - properties - Mozilla Firefox

File Edit View History Bookmarks Tools Help

redhat.comhttps://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&nodename=sayoko.salab.dfw

Most VisitedRed HatZimbraIntranetSalesForceSMESatellite.tc.redhat...Axiom TravelRed Hat Appscho.salab.dfw.redh...luci

Luci - cluster - nodes - proper...

Nodes

Add a Node

Configure

satoiko.salab.dfw.redhat.com

sakuko.salab.dfw.redhat.com

sakura.salab.dfw.redhat.com

sayoko.salab.dfw.redhat.com

sanoiko.salab.dfw.redhat.com

sakurako.salab.dfw.redhat.com

Services

Resources

Fallover Domains

Shared Fence Devices

Cluster daemons running on this node

| Daemon | Currently running | Enabled at start-up |
|-----------|-------------------|-------------------------------------|
| cman | yes | <input checked="" type="checkbox"/> |
| rgmanager | yes | <input checked="" type="checkbox"/> |

Update node daemon properties

Services on this Node

No cluster services are currently running here

Fallover Domain Membership

This node has no fallover domain membership

Main Fencing Method

Virtual Machine Fencing

Fence Type

Name

donU-virt-fence

Domain

guest1

Remove this instance

Remove this device

Add an instance

Add a fence device to this level

Update main fence properties

Backup Fencing Method

Add a fence device to this level

Update backup fence properties

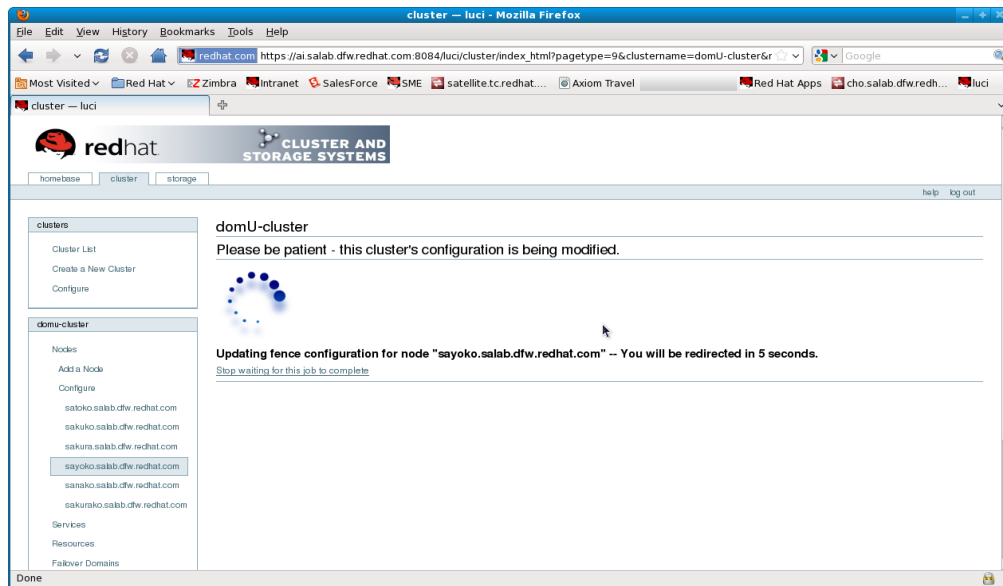
Done

127

Red Hat Summit 2009 | Thomas Cameron

RED HAT :: CHICAGO :: 2009
SUMMIT

The “Domain” name is the name of the virtual machine, in this case sayoko is running as guest1



Set up the clustered app on the domU cluster

Test by fencing a domU

Luci - cluster - nodes - properties - Mozilla Firefox

https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=9&nodename=satoko.salab.dfw

redhat CLUSTER AND STORAGE SYSTEMS

homebase cluster storage help log out

clusters

- Cluster List
- Create a New Cluster
- Configure

domU-cluster

Nodes

- Add a Node
- Configure
- satoko.salab.dfw.redhat.com
- sakuko.salab.dfw.redhat.com
- sakura.salab.dfw.redhat.com
- sayoko.salab.dfw.redhat.com
- sanako.salab.dfw.redhat.com
- sakurako.salab.dfw.redhat.com
- Services
- Resources
- Fallover Domains

Done

domU-cluster

Node Name: satoko.salab.dfw.redhat.com

Fence this node Go

Status: Cluster member

Show recent log activity for this node

Cluster daemons running on this node

| Daemon | Currently running | Enabled at start-up |
|-----------|-------------------|-------------------------------------|
| cman | yes | <input checked="" type="checkbox"/> |
| rgmanager | yes | <input checked="" type="checkbox"/> |

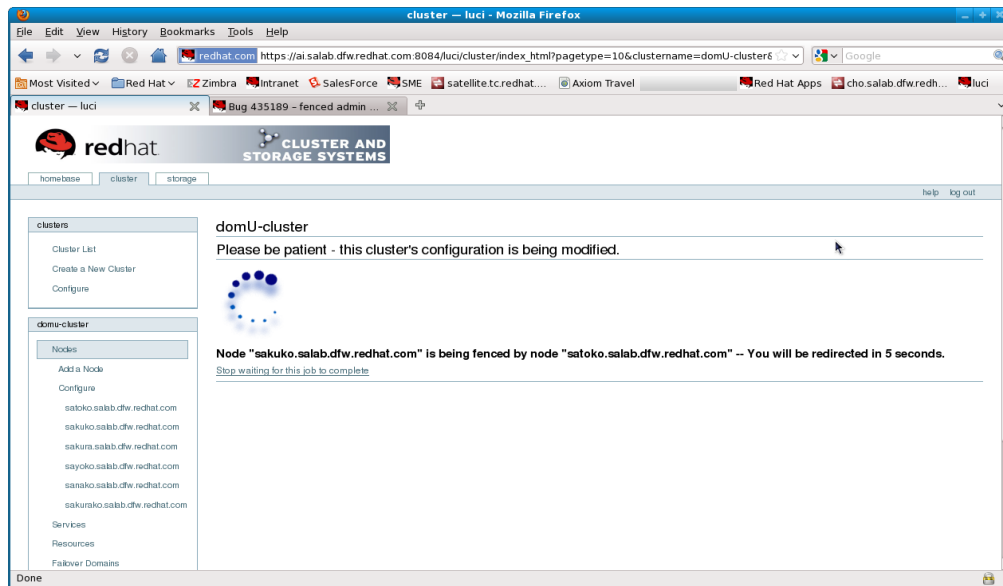
Update node daemon properties

Services on this Node

- No cluster services are currently running here

Fallover Domain Membership

- This node has no fallover domain membership



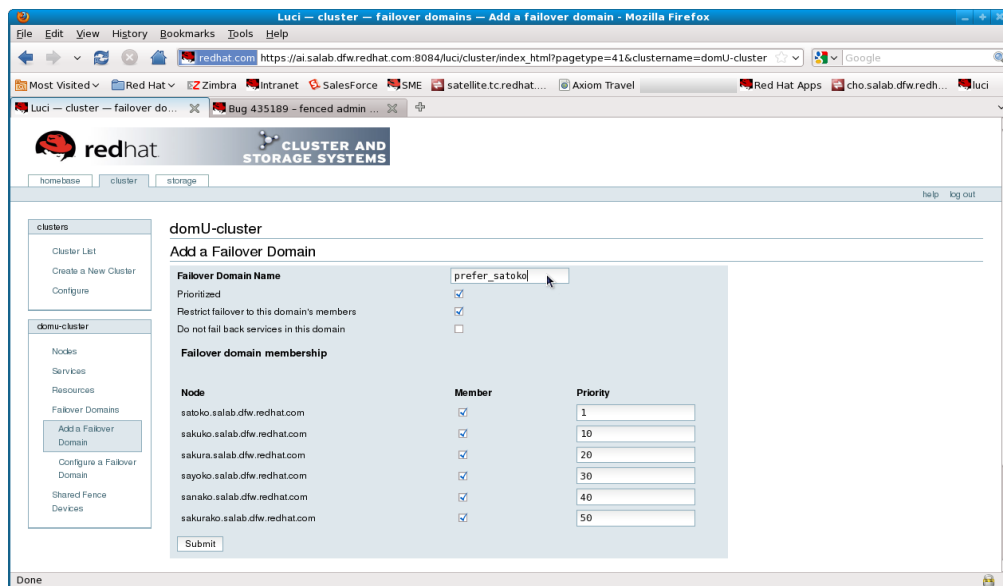
```
root@satoko:~  
File Edit View Terminal Help  
Aug 28 15:00:11 satoko fence_node[3646]: Fence of "sakuko.salab.dfw.redhat.com"  
was successful  
Aug 28 15:00:20 satoko openais[2203]: [TOTEM] entering GATHER state from 12.  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] entering GATHER state from 0.  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] Saving state aru 91 high seq recei  
ved 91  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] Storing new sequence id for ring 6  
c  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] entering COMMIT state.  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] entering RECOVERY state.  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] position [0] member 10.15.183.92:  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] previous ring seq 104 rep 10.15.18  
3.91  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] aru 91 high delivered 91 received  
flag 1  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] position [1] member 10.15.183.93:  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] previous ring seq 104 rep 10.15.18  
3.91  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] aru 91 high delivered 91 received  
flag 1  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] position [2] member 10.15.183.94:  
Aug 28 15:00:25 satoko openais[2203]: [TOTEM] previous ring seq 104 rep 10.15.18  
3.91
```

Set up the clustered app on the domU cluster

Define failover domain(s) for the app you want to run on this domU cluster

In this case, a simple Apache web server with an NFS mounted /var/www/html and a floating IP address

Here I've defined a FD prefer_satoko



Set up the clustered app on the domU cluster

Define resources for the service

- Cluster Tab

- Cluster Name

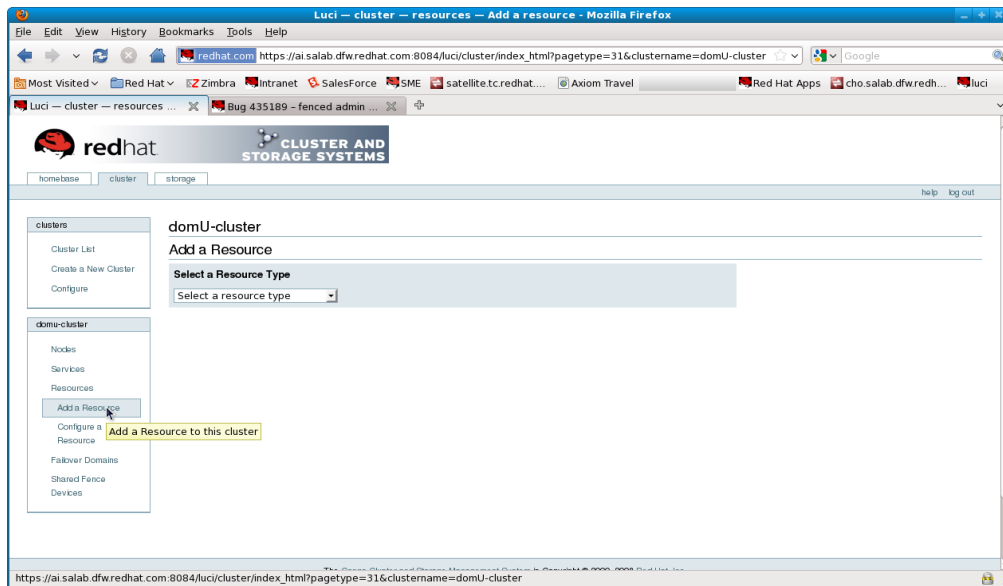
- Resources

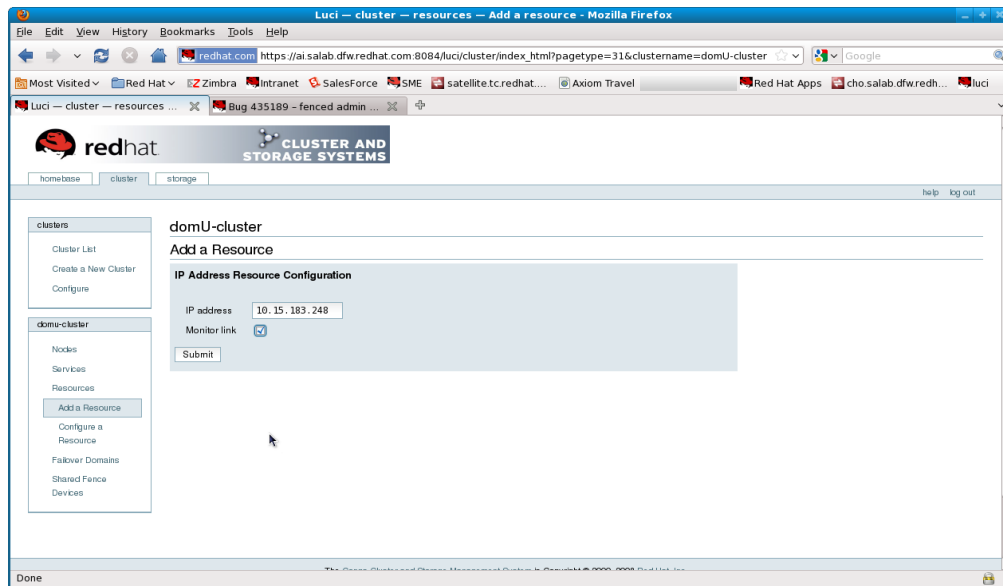
- Add a Resource

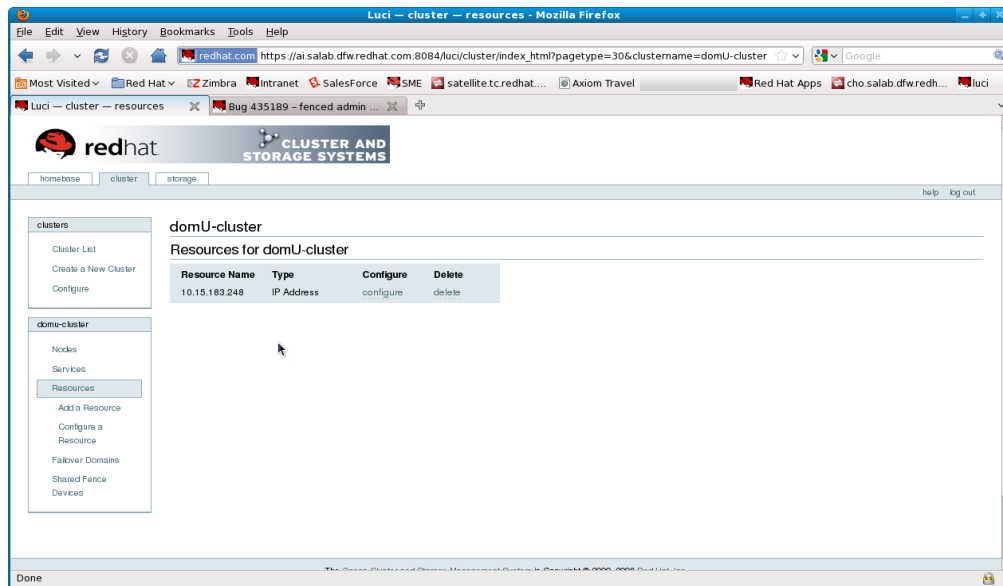
 - IP address

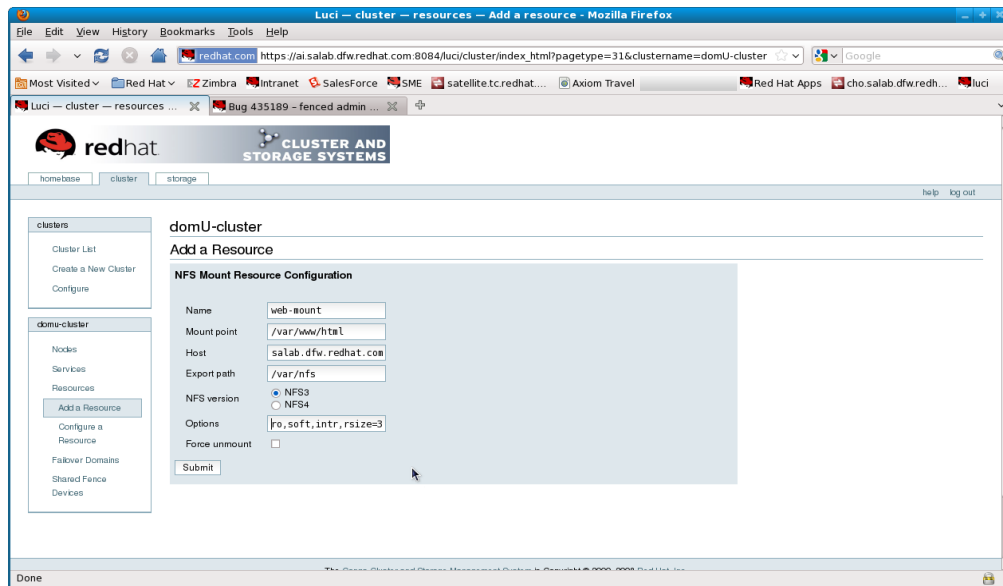
 - NFS mount

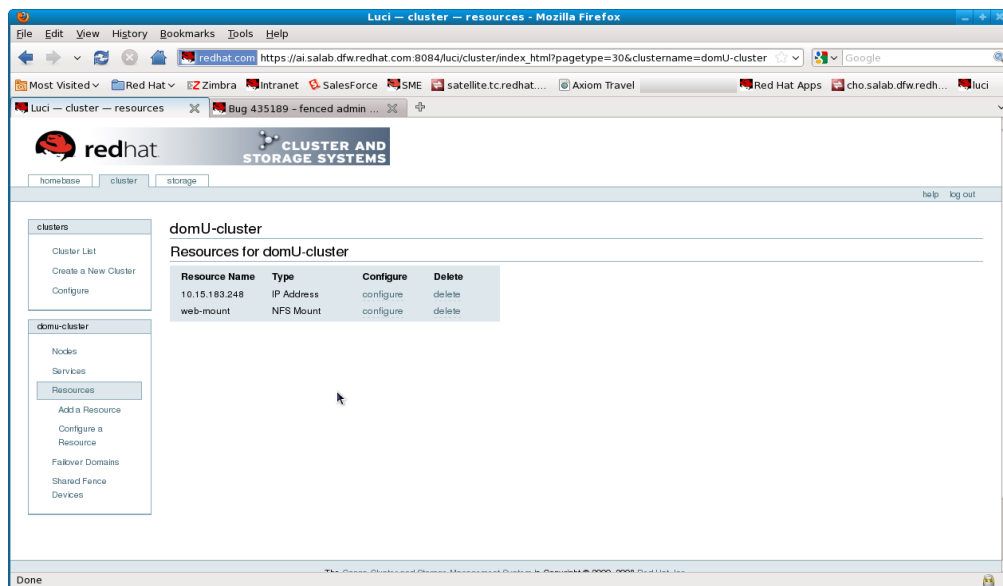
 - Script

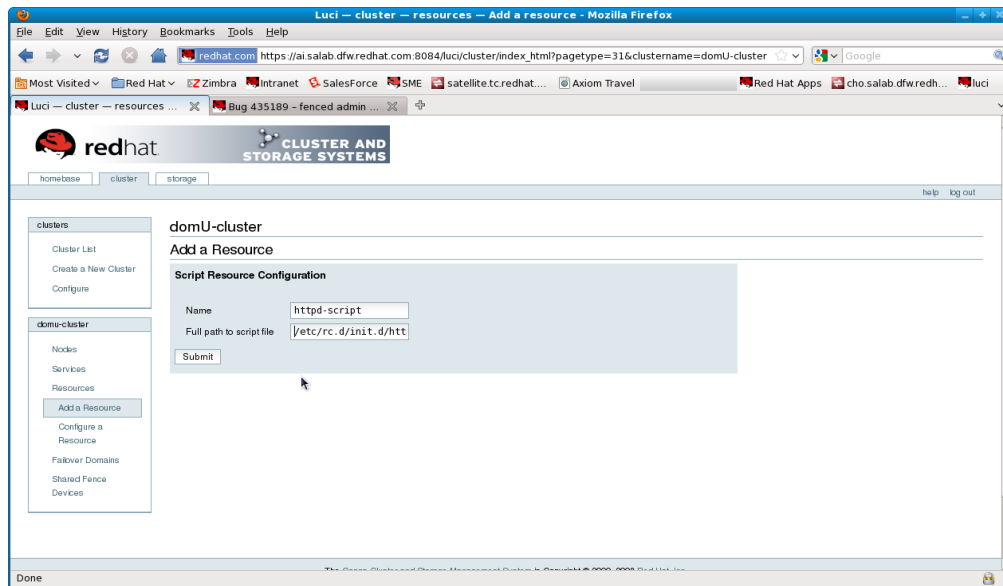


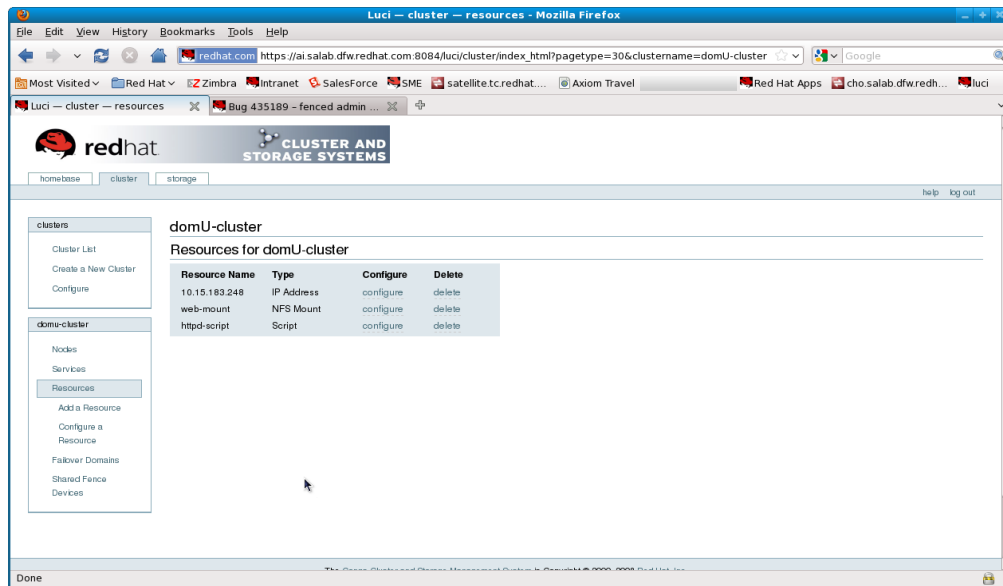












Set up the clustered app on the domU cluster

Define a service which uses all of those resources

Cluster Tab

Cluster Name

Services

Set up the clustered app on the domU cluster

Define a service which uses all of those resources

Add a Service

Name

Auto-start?

Enable NFS lock workarounds


Run exclusive

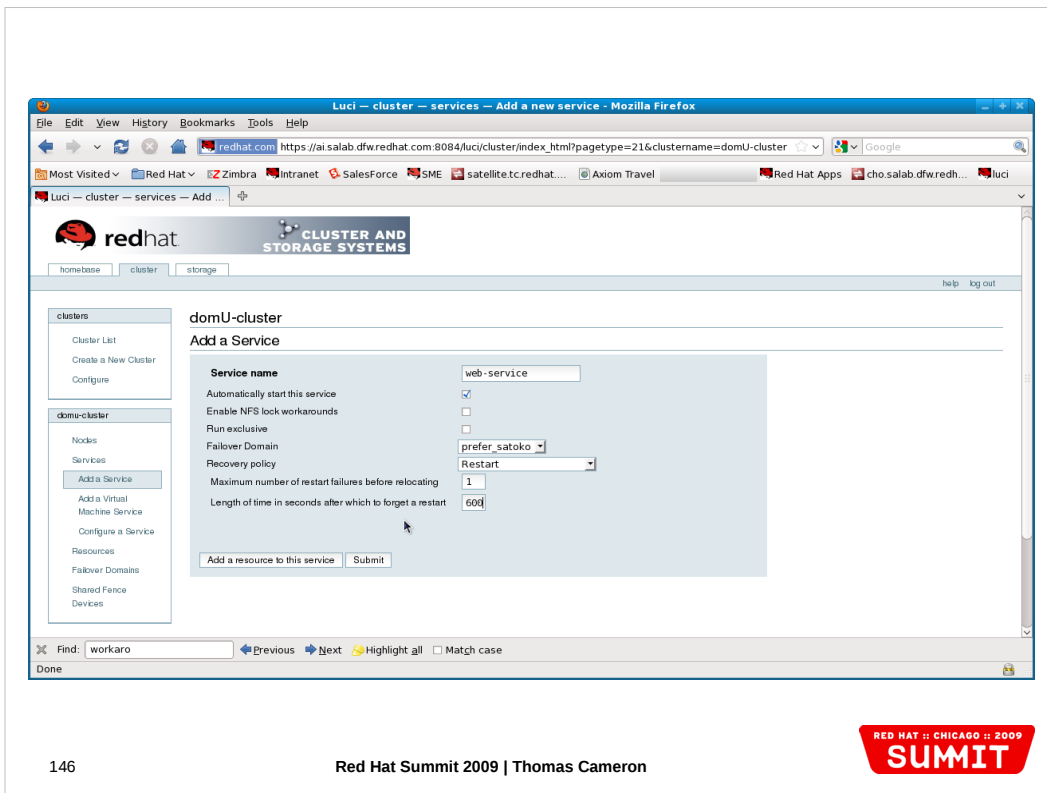
FD

Recovery Policy

Maximum number of restart failures before relocating

Length of time in seconds after which to forget a restart





Since this is not an HA callout like clunfslock (just failing over from one domU to another), there is no need to check the “Enable NFS lock workarounds” box.

Set up the clustered app on the domU cluster

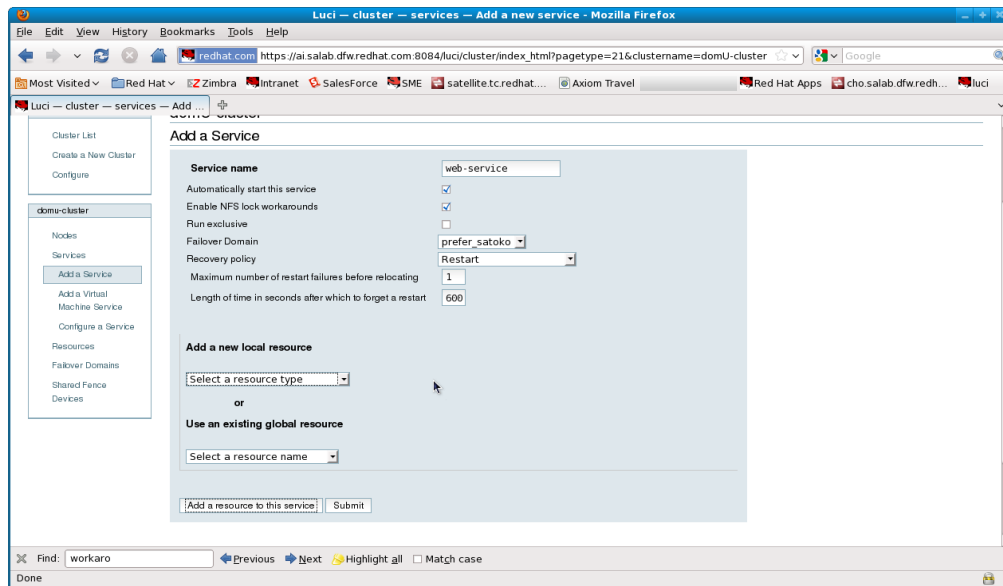
Define a service which uses all of those resources

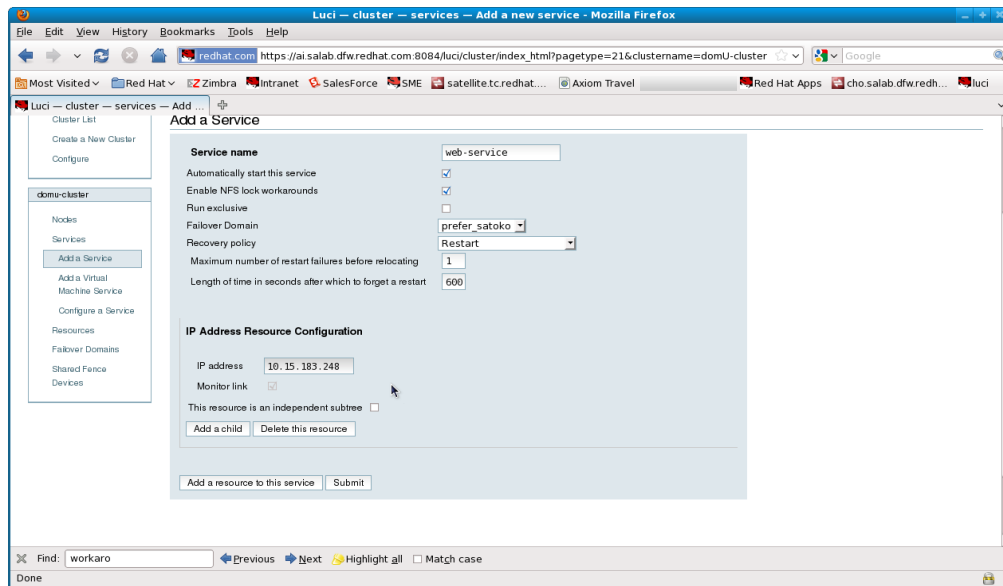
Add Global Resources to the service

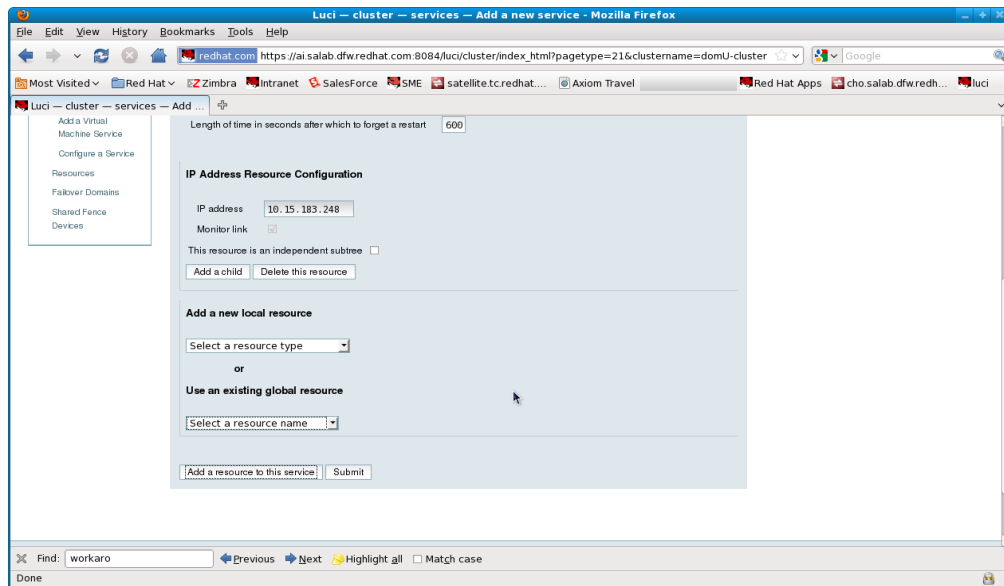
IP address

NFS Mount

Script







Luci — cluster — services — Add a new service - Mozilla Firefox

File Edit View History Bookmarks Tools Help

https://ai.salab.dfw.redhat.com:8084/luci/cluster/index_html?pagetype=21&clustername=domU-cluster

Most Visited Red Hat Zimbra Intranet Salesforce SME satellite.tc.redhat... Axiom Travel Red Hat Apps cho.salab.dfw.redh... luci

Luci — cluster — services — Add ...

Resources
Failover Domains
Shared Fence
Devices

IP Address Resource Configuration

IP address: 10.15.183.248

Monitor link: ☐

This resource is an independent subtree: ☐

Add a child Delete this resource

NFS Mount Resource Configuration

Name: web-mount

Mount point: /var/www/html

Host: ai.salab.dfw.redhat.

Export path: /var/nfs

NFS version: ☒ NFS3 ☐ NFS4

Options: ro,soft,intr,rsize=3

Force unmount: ☐

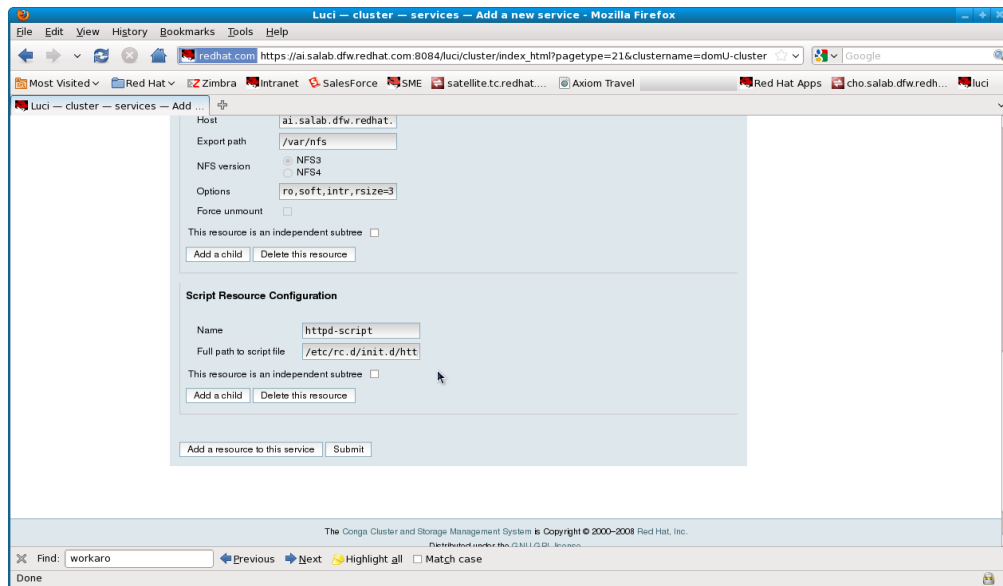
This resource is an independent subtree: ☐

Add a child Delete this resource

Add a resource to this service Submit

Find: workaro Previous Next Highlight all Match case

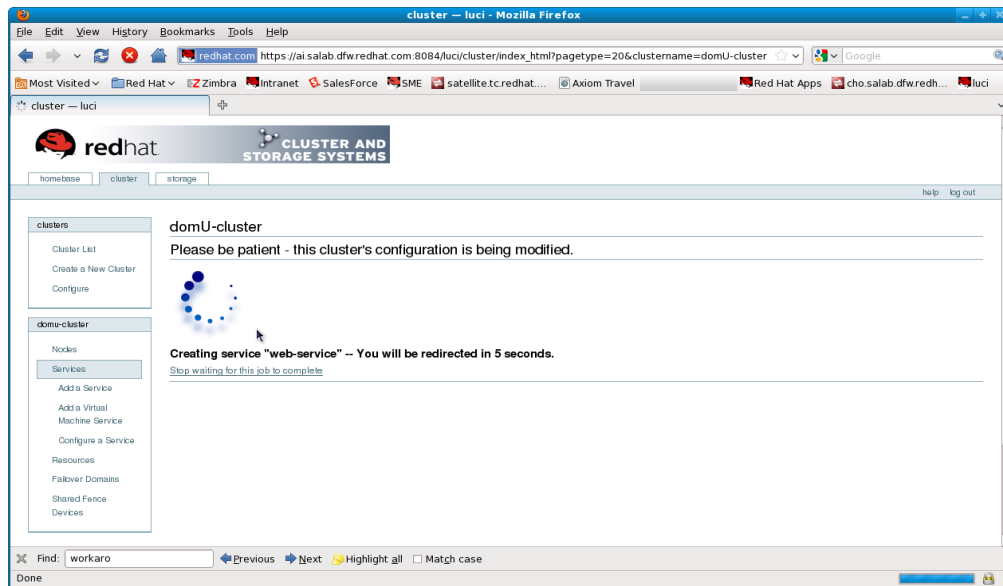
Done



Set up the clustered app on the domU cluster

Submit and confirm

Open the IP address defined in the resource definition in your web browser




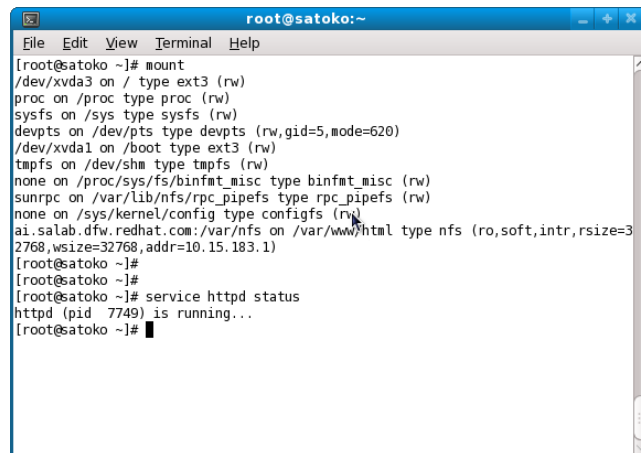
155

Red Hat Summit 2009 | Thomas Cameron

RED HAT :: CHICAGO :: 2009
SUMMIT

Remember that it may take a few seconds for the service to go green. If it comes up red, just refresh the screen a few times.

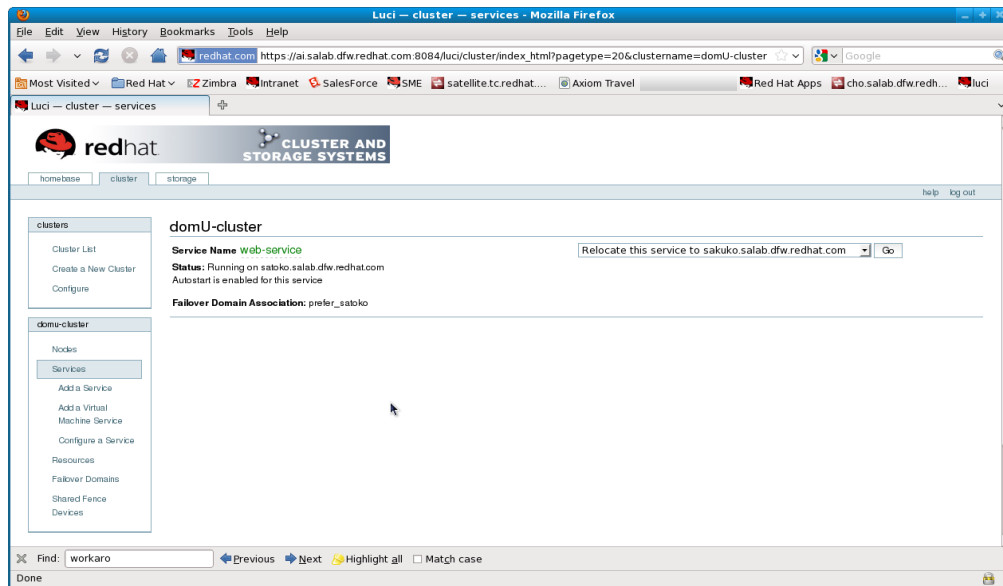


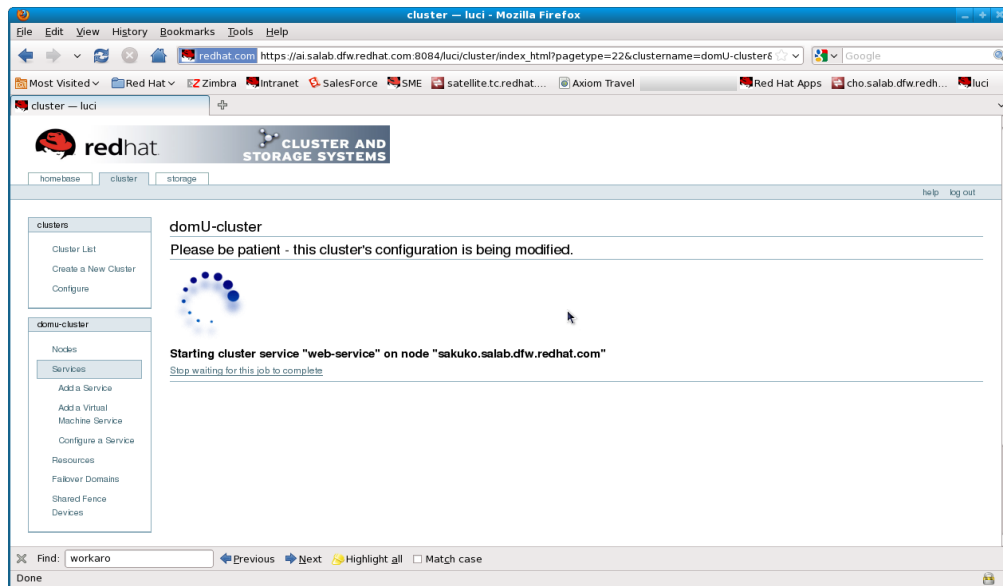


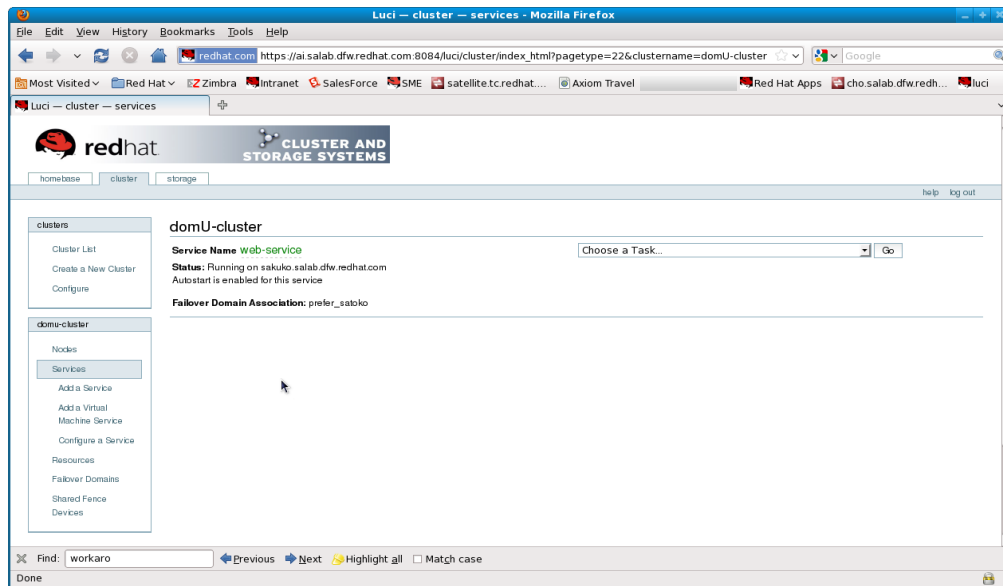
```
root@satoko: ~  
File Edit View Terminal Help  
[root@satoko ~]# mount  
/dev/xvda3 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/xvda1 on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /sys/kernel/config type configfs (rw)  
ai.salab.dfw.redhat.com:/var/nfs on /var/www/html type nfs (ro,soft,intr,rsize=32768,wsiz  
e=32768,addr=10.15.183.1)  
[root@satoko ~]#  
[root@satoko ~]#  
[root@satoko ~]# service httpd status  
httpd (pid 7749) is running...  
[root@satoko ~]#
```

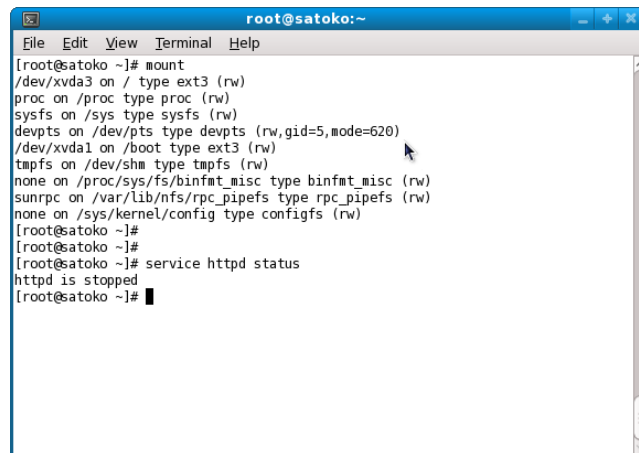
Set up the clustered app on the domU cluster

Test migration and failover









```
root@satoko:~  
File Edit View Terminal Help  
[root@satoko ~]# mount  
/dev/xvda3 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/xvda1 on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /sys/kernel/config type configfs (rw)  
[root@satoko ~]#  
[root@satoko ~]#  
[root@satoko ~]# service httpd status  
httpd is stopped  
[root@satoko ~]#
```



```
root@sakuko:~  
File Edit View Terminal Help  
Connection to satoko.salab.dfw.redhat.com closed.  
[root@ai ~]# ssh sakuko.salab.dfw.redhat.com  
Last login: Fri Aug 28 14:58:25 2009 from ai.salab.dfw.redhat.com  
[root@sakuko ~]# mount  
/dev/xvda3 on / type ext3 (rw)  
proc on /proc type proc (rw)  
sysfs on /sys type sysfs (rw)  
devpts on /dev/pts type devpts (rw,gid=5,mode=620)  
/dev/xvda1 on /boot type ext3 (rw)  
tmpfs on /dev/shm type tmpfs (rw)  
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)  
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)  
none on /sys/kernel/config type configfs (rw)  
ai.salab.dfw.redhat.com:/var/nfs on /var/nfs type nfs (ro,soft,intr,rsize=32768,wsiz=32768,addr=10.15.183.1)  
[root@sakuko ~]#  
[root@sakuko ~]#  
[root@sakuko ~]# service httpd status  
httpd (pid 7634) is running...  
[root@sakuko ~]#
```

```
root@sakuko:~  
File Edit View Terminal Help  
[root@sakuko ~]# clustat  
Cluster Status for domU-cluster @ Fri Aug 28 16:12:17 2009  
Member Status: Quorate  
  
Member Name                                ID  Status  
-----  
satoko.salab.dfw.redhat.com                1  Online, rgmanager  
sakuko.salab.dfw.redhat.com                2  Online, Local, rgmanager  
sakura.salab.dfw.redhat.com                3  Online, rgmanager  
sayoko.salab.dfw.redhat.com                4  Online, rgmanager  
sanako.salab.dfw.redhat.com                5  Online, rgmanager  
sakurako.salab.dfw.redhat.com              6  Online, rgmanager  
  
Service Name      Owner (Last)      State  
-----  
service:web-service  sakuko.salab.dfw.redhat.com  started  
[root@sakuko ~]#
```

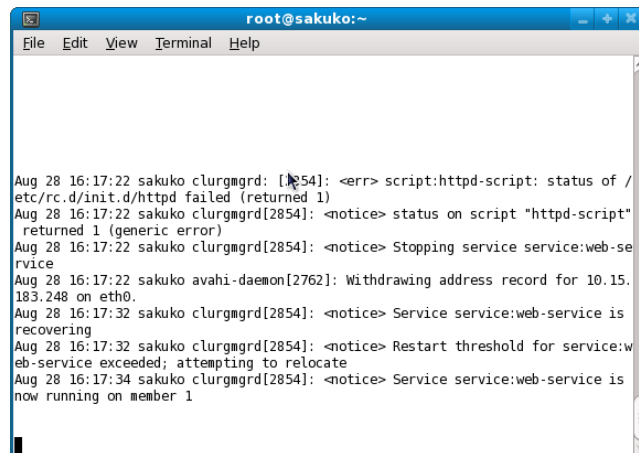
Set up the clustered app on the domU cluster

Test, test, test

Kill the app

Crash the domU

Crash the dom0



```
root@sakuko:~  
File Edit View Terminal Help  
  
Aug 28 16:17:22 sakuko clurgmgrd[18254]: <err> script:httpd-script: status of /  
etc/rc.d/init.d/httpd failed (returned 1)  
Aug 28 16:17:22 sakuko clurgmgrd[2854]: <notice> status on script "httpd-script"  
returned 1 (generic error)  
Aug 28 16:17:22 sakuko clurgmgrd[2854]: <notice> Stopping service service:web-se  
rvice  
Aug 28 16:17:22 sakuko avahi-daemon[2762]: Withdrawing address record for 10.15.  
183.248 on eth0.  
Aug 28 16:17:32 sakuko clurgmgrd[2854]: <notice> Service service:web-service is  
recovering  
Aug 28 16:17:32 sakuko clurgmgrd[2854]: <notice> Restart threshold for service:w  
eb-service exceeded; attempting to relocate  
Aug 28 16:17:34 sakuko clurgmgrd[2854]: <notice> Service service:web-service is  
now running on member 1
```

In this slide I've issued "pkill -9 httpd" twice. This demonstrates that the cluster has restarted the service once, then relocated per our FD thresholds.

WOW! You made it this far!

This should give you a good start - it is not, and can not be, comprehensive.

To learn a lot more about virtualization, take RH401:

<http://tinyurl.com/5gnj2f>

To learn a lot more about clustering, look at RH436:

<http://tinyurl.com/o89lnr>

If you need help getting this all set up, contact Red Hat Consulting:

<http://www.redhat.com/consulting/>

Thank you very much!

Feel free to reach out to me at thomas@redhat.com, or Lon
at lh@redhat.com

Updated slides at

<http://people.redhat.com/tcameron>

<http://www.redhat.com/promo/summit>

Please take the survey at:

<http://www.redhat.com/summit-survey>

QUESTIONS?

**TELL US WHAT YOU THINK:
[REDHAT.COM/SUMMIT-SURVEY](https://redhat.com/summit-survey)**