



# RED HAT NETWORK SATELLITE HIGH AVAILABILITY GUIDE

BY WAYMAN SMITH, RED HAT CONSULTANT

---

Red Hat® Network (RHN) Satellite Server is an easy-to-use systems management platform for your growing Red Hat Enterprise Linux® infrastructure. Built on open standards, RHN Satellite provides powerful systems administration capabilities such as management, provisioning and monitoring for large deployments. Satellite allows you to manage many servers as easily as you would one.

Red Hat Cluster Suite allows applications or services to be deployed in high availability configurations so that they are always operational—bringing “scale-out” capabilities to enterprise Linux deployments.

This paper describes the procedure for combining these two technologies to provide high-availability Satellite Server to your environment.



## TABLE OF CONTENTS

---

### 3 INTRODUCTION

- 3 About Red Hat Cluster Suite
- 3 About Red Hat Network Satellite

### 4 ENVIRONMENT DETAILS

- 4 Hardware
- 4 Software

### 4 PROCEDURES

- 4 Environment Preparation
  - 4 Hostname Resolution
  - 6 Add virtual IP address
  - 7 Create LVM partitions
- 13 Red Hat Satellite Network Server
  - 13 Satellite installation (node 1)
  - 22 Disable Automatic Satellite Startup
  - 24 Satellite Installation (node 2)
  - 29 Verify Satellite operates properly
- 30 Red Hat Cluster Suite
  - 30 Red Hat Cluster Suite subscription
- 35 Red Hat Cluster Suite Installation
  - 37 Configure quorum disk
  - 39 Configure cluster nodes
  - 44 Add Fencing Device
  - 45 Add Fence Devices to node
- 46 Configure cluster properties
  - 46 Configuring failover domains
  - 47 Configure cluster resources
- 50 Activate Cluster suite
- 53 Test cluster failover ability

### 55 RECOMMENDATIONS

### 55 CONCLUSION

- 55 Red Hat Enterprise Linux
- 55 Red Hat Network Satellite
- 56 Red Hat Cluster Suite **Introduction**



## INTRODUCTION

---

Combining the high availability features of Red Hat Cluster Suite with the system management capability of Red Hat Network Satellite allows you effectively manage your client systems with the assurance that the application will remain active and functioning properly. This intermingling of Red Hat product offerings is possible because of the Red Hat Enterprise Linux operating system. Red Hat Enterprise Linux serves as the base platform for which these applications are developed, allowing Red Hat Cluster Suite and Red Hat Satellite Server to function seamlessly with Red Hat Enterprise Linux. It is the ability to combine products, as demonstrated in this paper, that allows customers to utilize Red Hat Enterprise Linux and other product offerings in a unique and fully functional manner.

### ABOUT RED HAT CLUSTER SUITE

For applications that require maximum up-time and high availability, Red Hat accommodates with Red Hat Cluster Suite. With Red Hat Cluster Suite, both standard services (Apache) and custom applications can be configured to operate in a fail-over configuration that maximizes application up-time and can effectively migrate applications in the event of hardware failure.

Red Hat Cluster Suite can combine up to 128 systems together in a cluster configuration to provide one or several applications to clients. To provide a consistent view to client systems, clusters can share resources (IP address, mount point, etc.) to all nodes participating to the cluster.

### ABOUT RED HAT NETWORK SATELLITE

Red Hat Network Satellite provides effective system management of large enterprise environments with an intuitive, usable, and customizable interface. Normal system administration tasks that require physical system interaction (logging in, executing command, verifying result, etc.) can be consolidated into one web interface. Not only can system administration tasks be performed from one interface, but by grouping systems together, tasks can be performed on a number of systems at once.

The management module provides the user with the ability to organize systems in static and dynamic groups to perform a wide range of tasks. In addition to system groups, the administrator can create software channels that supply packages, updates, and errata to all systems subscribed. Software channels can even be created with the purpose of software lifecycle management by creating development channels with the latest package updates, testing channels where only specific updates are imported, and production channels in which the administrator has total control of the packages present in that channel.

The provisioning module allows for extended system management beyond the package level. With provisioning, the system administrator has the ability to create system snapshots, software profile lists, kickstarts, and configuration management channels. Systems can now compare, duplicate, and re-provision all from one common interface. Changes to a system can be rolled back with ease through system snapshots.

The monitoring module provides effective monitoring of systems and applications all from within the Satellite web interface. By deploying monitoring probes, system performance can now be tracked, and administrators can be alerted of the latest changes in system status. In addition to standard probes for temperature, storage space, etc., administrators can also deploy custom probes to meet their specific needs.

What makes Red Hat Network so beneficial is that the functionality of these management features is built into every copy of Red Hat Enterprise Linux, whereas other third-party applications require installation and configuration that can be lengthy, complicated, and involved. Environment details



## ENVIRONMENT DETAILS

---

The environment in which this procedure was performed was as follows:

### HARDWARE

Two IBM x3650 servers

- Quad x5450 3GHz
- 16GB
- Diskless - Boot from SAN
- Dual 1000Mbps onboard Ethernet
- Quad 1000Mbps adapter Ethernet
- Two dual channel 5Gn fiber channel attached to SAN
- RSAII

One IBM DS3400 Disk Array

- Redundant dual ported RAID controllers
- Redundant power supply
- Six 146GB disks

Fencing Device: IBM RSA II

### SOFTWARE

- Red Hat Enterprise Linux 5.2 x86\_64
- Red Hat Network Satellite 5.1 x86\_64
- Red Hat Cluster Suite 4.6

## PROCEDURES

---

### ENVIRONMENT PREPARATION

Red Hat Cluster Suite (RHCS) nodes will share resources (IP Address/hostname, Fibre-channel storage), so it is necessary to configure these resources prior to Red Hat Cluster Satellite installation.

#### Hostname resolution

Since a single IP address and hostname will be shared between both nodes in the cluster, it is necessary to have proper hostname resolution in place for that IP address/hostname pair. In this procedure proper resolution was achieved through entries in the `/etc/hosts` file. Additional methods of resolution can also be used (DNS).





FIG 1: HOSTS FILE

```
root@localhost:~  
File Edit View Terminal Tabs Help  
[root@node1 ~]# cat /etc/hosts  
# Do not remove the following line, or various programs  
# that require network functionality will fail.  
127.0.0.1          localhost.localdomain localhost  
::1               localhost6.localdomain6 localhost6  
192.168.1.30      sat.example.com sat  
192.168.1.10      node1.example.com node1  
192.168.1.20      node2.example.com node1  
[root@node1 ~]# hostname  
node1.example.com  
[root@node1 ~]#
```



FIG 2: VIRTUAL IP ADDRESS

```
root@localhost:/etc/sysconfig/network-scripts
File Edit View Terminal Tabs Help
[root@node1 network-scripts]# ip addr add 192.168.1.30/24 dev bond0
[root@node1 network-scripts]# ip addr list dev bond0
4: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether 00:17:de:36:26:01 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.10/24 brd 192.168.1.255 scope global bond0
    inet 192.168.1.30/24 scope global secondary bond0
    inet6 fe80::217:deff:fe36:2601/64 scope link
        valid_lft forever preferred_lft forever
[root@node1 network-scripts]# ip addr del 192.168.1.30/24 dev bond0
[root@node1 network-scripts]# ip addr list dev bond0
4: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether 00:17:de:36:26:01 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.10/24 brd 192.168.1.255 scope global bond0
    inet6 fe80::217:deff:fe36:2601/64 scope link
        valid_lft forever preferred_lft forever
[root@node1 network-scripts]#
```

### Add virtual IP address

It is necessary to manually add virtual IP address when testing Satellite functionality on each individual node prior to cluster implementation.

To configure virtual IP address, use the following command:

```
# ip addr add 192.168.1.30/24 dev bond0
```

Verify virtual address has been created:

```
#ip addr list dev bond0
bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 100
    link/ether 00:15:58:c4:17:f8 brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.216/24 brd 192.168.0.255 scope global eth0
    inet 192.168.1.30/24 scope global secondary eth0
```

To remove virtual IP address, use the following command:

```
# ip addr del 192.168.1.30/24 dev bond0
```



FIG 3: FDISK OUTPUT

```
root@node1:~  
File Edit View Terminal Tabs Help  
[root@node1 ~]# fdisk -l  
  
Disk /dev/hda: 20.9 GB, 20971520000 bytes  
255 heads, 63 sectors/track, 2549 cylinders  
Units = cylinders of 16065 * 512 = 8225280 bytes  
  
   Device Boot      Start         End      Blocks   Id  System  
/dev/hda1 *          1           13       104391   83  Linux  
/dev/hda2             14        2549      20370420  8e  Linux LVM  
  
Disk /dev/sda: 62.9 GB, 62914560000 bytes  
64 heads, 32 sectors/track, 60000 cylinders  
Units = cylinders of 2048 * 512 = 1048576 bytes  
  
Disk /dev/sda doesn't contain a valid partition table  
[root@node1 ~]#
```

### Create LVM partitions

To create LVM partitions, follow the procedures below on node 1 system  
Verify storage can be seen by cluster node:

```
# fdisk -l  
# fdisk /dev/sda
```

Create physical partitions:

```
# fdisk /dev/sda
```

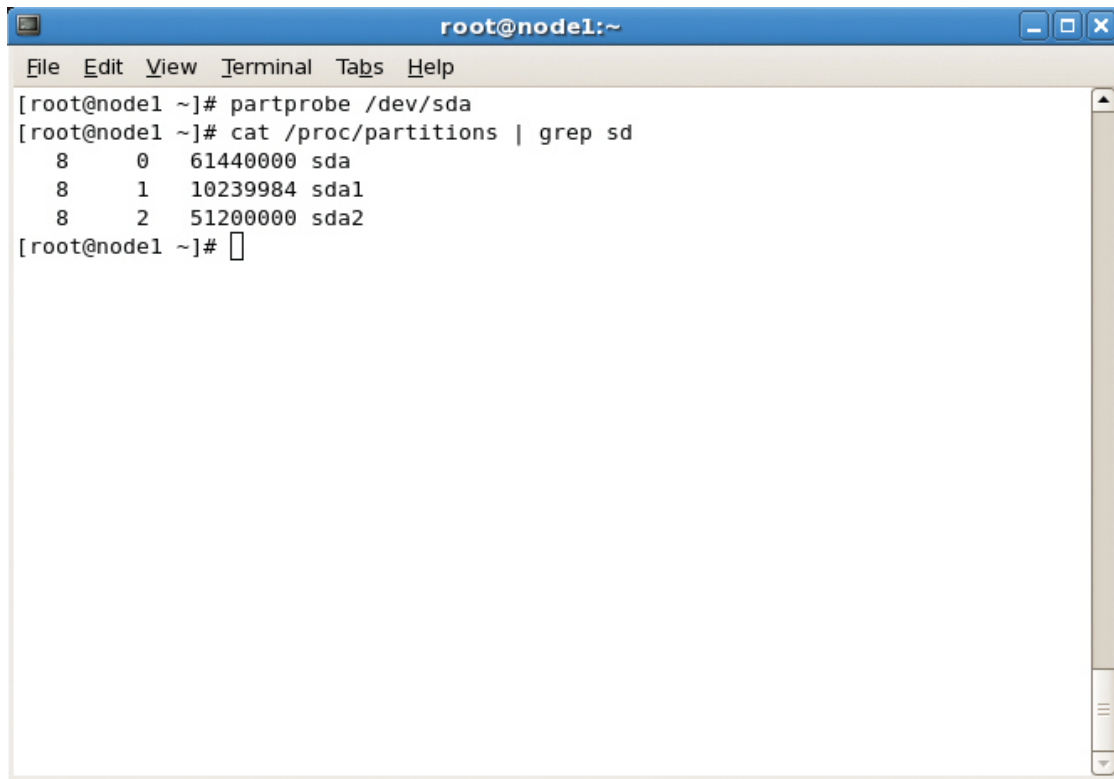
Partition size is dependent on the amount of software channels (custom and base) that will be provided by satellite.

Once partition(s) have been created, verify existence and create lvm physical volumes:

```
# partprobe /dev/sda  
# cat /proc/partitions
```



FIG 4: PARTPROBE OUTPUT



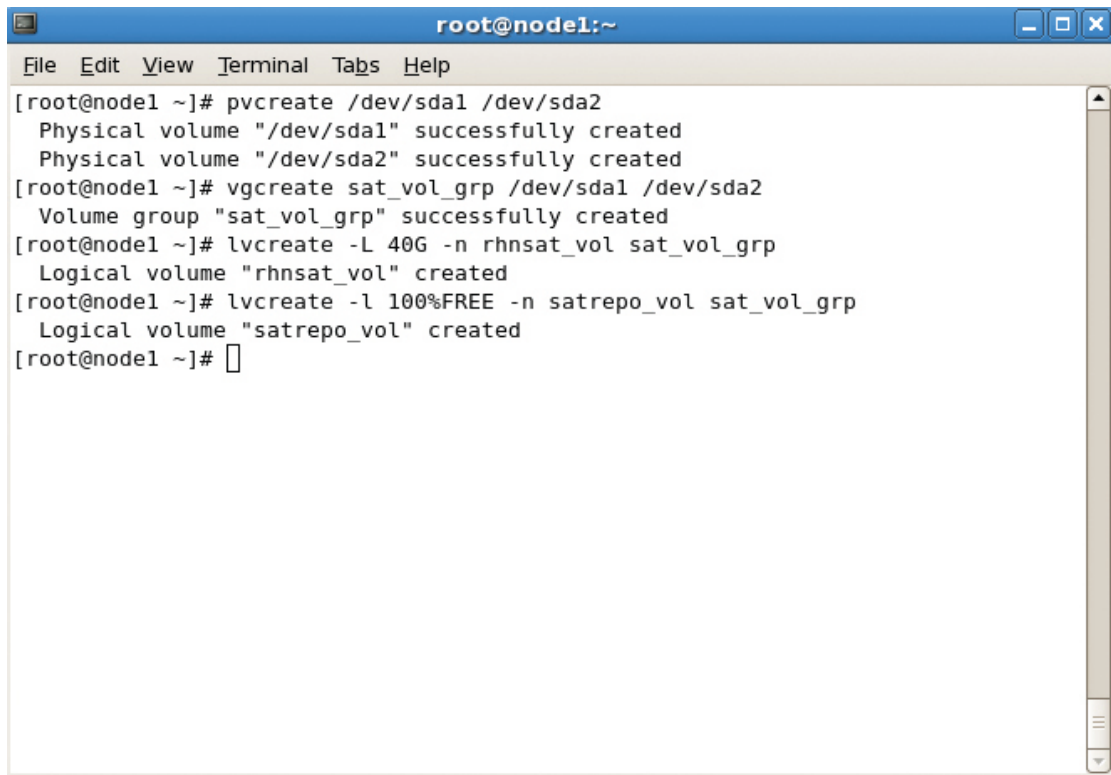
```
root@node1:~  
File Edit View Terminal Tabs Help  
[root@node1 ~]# partprobe /dev/sda  
[root@node1 ~]# cat /proc/partitions | grep sd  
 8    0  61440000 sda  
 8    1  10239984 sda1  
 8    2  51200000 sda2  
[root@node1 ~]#
```

At this point LVM physical partitions can be created:

```
# pvcreate /dev/sdc1 /dev/sdc2  
# vgcreate -c y sat_vol_grp /dev/sdc1 /dev/sdc2  
# lvcreate -L 40G -n rhnsat_vol sat_vol_grp  
# lvcreate -l 100%FREE -n satrepo_vol sat_vol_grp
```



FIG 5: LVM OUTPUT

A terminal window titled "root@node1:~" with a menu bar containing "File", "Edit", "View", "Terminal", "Tabs", and "Help". The terminal displays the following commands and their outputs:

```
[root@node1 ~]# pvcreate /dev/sda1 /dev/sda2
  Physical volume "/dev/sda1" successfully created
  Physical volume "/dev/sda2" successfully created
[root@node1 ~]# vgcreate sat_vol_grp /dev/sda1 /dev/sda2
  Volume group "sat_vol_grp" successfully created
[root@node1 ~]# lvcreate -L 40G -n rhnsat_vol sat_vol_grp
  Logical volume "rhnsat_vol" created
[root@node1 ~]# lvcreate -l 100%FREE -n satrepo_vol sat_vol_grp
  Logical volume "satrepo_vol" created
[root@node1 ~]#
```



FIG 6: LV DISPLAY OUTPUT

```
root@node1:~  
File Edit View Terminal Tabs Help  
LV Name          /dev/sat_vol_grp/rhnsat_vol  
VG Name          sat_vol_grp  
LV UUID          6j68rE-nUG9-puwa-Z6bd-rago-0I7L-WKgBmk  
LV Write Access  read/write  
LV Status        available  
# open          0  
LV Size          40.00 GB  
Current LE       10240  
Segments         1  
Allocation       inherit  
Read ahead sectors auto  
- currently set to 256  
Block device     253:2  
  
--- Logical volume ---  
LV Name          /dev/sat_vol_grp/satrepo_vol  
VG Name          sat_vol_grp  
LV UUID          thZ40u-PhXi-9H7a-Akby-0u7p-xvvy-EFV1e5  
LV Write Access  read/write  
LV Status        available  
# open          0  
LV Size          18.59 GB  
Current LE       4758  
Segments         2  
Allocation       inherit  
Read ahead sectors auto  
- currently set to 256  
Block device     253:3
```

It is now necessary to edit `lvm.conf (/etc/lvm/lvm.conf)` and change `locking_type = 3` to enable built-in cluster locking.



FIG 7: LVM.CONF LOCKING\_TYPE

A terminal window titled "root@node1:~" with a menu bar containing "File", "Edit", "View", "Terminal", "Tabs", and "Help". The terminal shows the command `grep -B 4 locking_type /etc/lvm/lvm.conf` and its output. The output consists of several lines of configuration comments and a setting: `locking_type = 3`. The terminal prompt is `[root@node1 ~]#` with a cursor.

```
root@node1:~  
File Edit View Terminal Tabs Help  
[root@node1 ~]# grep -B 4 locking_type /etc/lvm/lvm.conf  
# Turn locking off by setting to 0 (dangerous: risks metadata corruption  
# if LVM2 commands get run concurrently).  
# Type 2 uses the external shared library locking_library.  
# Type 3 uses built-in clustered locking.  
locking_type = 3  
--  
# Search this directory first for shared libraries.  
#  library_dir = "/lib"  
  
# The external locking library to load if locking_type is set to 2.  
[root@node1 ~]#
```

Create ext3 filesystem for logical volumes:

```
# mkfs.ext3 /dev/sat_vol_group/rhnsat_vol
```



FIG 8: EXT3 FORMAT RED HAT NETWORK SATELLITE STORAGE

```
root@node1:~  
File Edit View Terminal Tabs Help  
[root@node1 ~]# mkfs.ext3 /dev/sat_vol_grp/rhnsat_vol  
mke2fs 1.39 (29-May-2006)  
Filesystem label=  
OS type: Linux  
Block size=4096 (log=2)  
Fragment size=4096 (log=2)  
5242880 inodes, 10485760 blocks  
524288 blocks (5.00%) reserved for the super user  
First data block=0  
Maximum filesystem blocks=0  
320 block groups  
32768 blocks per group, 32768 fragments per group  
16384 inodes per group  
Superblock backups stored on blocks:  
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,  
    4096000, 7962624  
  
Writing inode tables: done  
Creating journal (32768 blocks): done  
Writing superblocks and filesystem accounting information: done  
  
This filesystem will be automatically checked every 32 mounts or  
180 days, whichever comes first. Use tune2fs -c or -i to override.  
[root@node1 ~]#
```

```
# mkfs.ext3 /dev/sat_vol_group/satrepo_vol
```





FIG 9: EXT3 FORMAT SATELLITE STORAGE

```
root@node1:~  
File Edit View Terminal Tabs Help  
[root@node1 ~]# mkfs.ext3 /dev/sat_vol_grp/satrepo_vol  
mke2fs 1.39 (29-May-2006)  
Filesystem label=  
OS type: Linux  
Block size=4096 (log=2)  
Fragment size=4096 (log=2)  
2436448 inodes, 4872192 blocks  
243609 blocks (5.00%) reserved for the super user  
First data block=0  
Maximum filesystem blocks=0  
149 block groups  
32768 blocks per group, 32768 fragments per group  
16352 inodes per group  
Superblock backups stored on blocks:  
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,  
    4096000  
  
Writing inode tables: done  
Creating journal (32768 blocks): done  
Writing superblocks and filesystem accounting information: done  
  
This filesystem will be automatically checked every 36 mounts or  
180 days, whichever comes first. Use tune2fs -c or -i to override.  
[root@node1 ~]#
```

## RED HAT NETWORK SATELLITE SERVER

### Satellite installation (Node 1)

IP address, hostname, and mount point shared resources must be in place on Node 1 prior to the Red Hat Network Satellite Server installation process beginning. Please refer to the above environment preparation procedures.

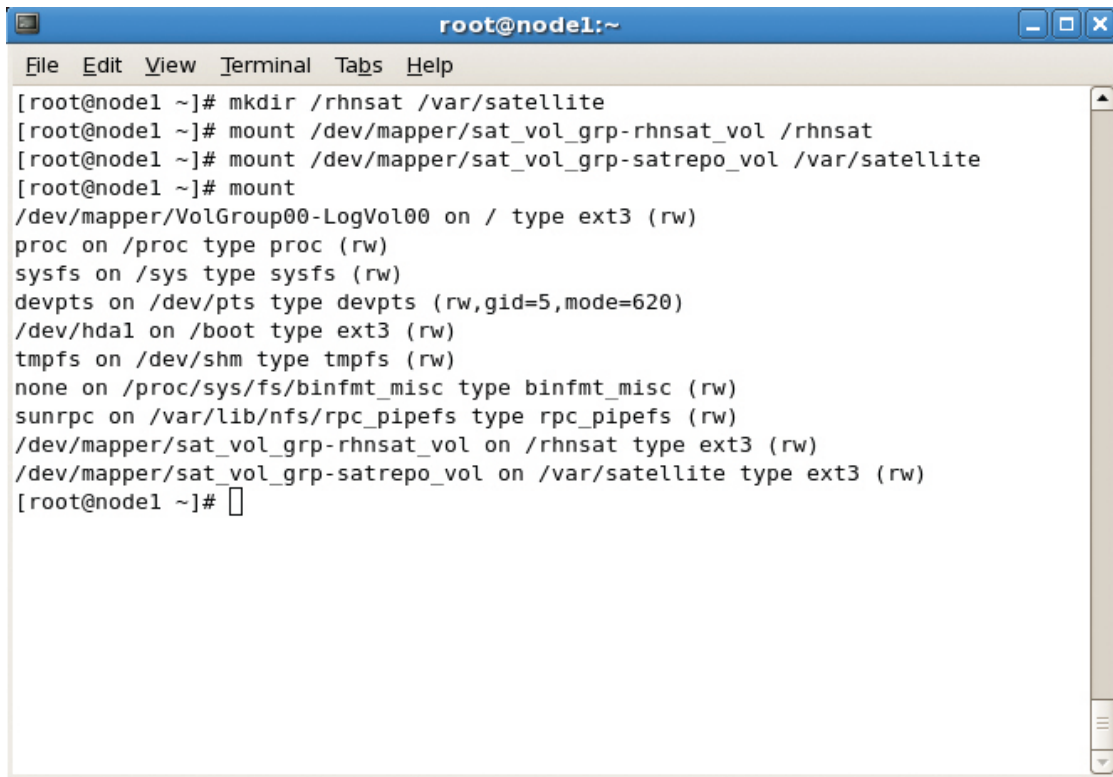
Preparing cluster node

1. Mount necessary storage devices.

```
# mkdir /rhnsat; mount /dev/mapper/sat_vol_grp-rhnsat_vol /rhnsat  
# mkdir /var/satellite; mount /dev/mapper/sat_vol_grp-satrepo_vol /var/satellite
```



FIG 10: MOUNT SHARED DRIVE



```
root@node1:~  
File Edit View Terminal Tabs Help  
[root@node1 ~]# mkdir /rhnsat /var/satellite  
[root@node1 ~]# mount /dev/mapper/sat_vol_grp-rhnsat_vol /rhnsat  
[root@node1 ~]# mount /dev/mapper/sat_vol_grp-satrepo_vol /var/satellite  
[root@node1 ~]# 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/hda1 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)  
/dev/mapper/sat_vol_grp-rhnsat_vol on /rhnsat type ext3 (rw)  
/dev/mapper/sat_vol_grp-satrepo_vol on /var/satellite type ext3 (rw)  
[root@node1 ~]#
```

2. Verify root user has write permissions on mounted directories.

```
# touch /var/satellite/file1; touch /rhnsat/file1
```

#### Satellite Installation

1. Insert and mount Satellite installation media.

```
# mount -o loop /tmp/Satellite-5.2-RHEL5.iso /mnt
```



FIG 11: MOUNT SAT ISO FILE

A terminal window titled "root@node1:~" with a menu bar containing "File", "Edit", "View", "Terminal", "Tabs", and "Help". The terminal shows the following commands and output:


```
[root@node1 ~]# mount -o loop /tmp/Satellite-5.2.0-RHEL5-re20081028.3-i386-embedded-oracle.iso /mnt
[root@node1 ~]# ls /mnt
EmbeddedDB  install  install.pl  README  Satellite  TRANS.TBL  updates
[root@node1 ~]#
```

2. Begin installation.

```
# ./install.pl --disconnected -answer-file=/tmp/answers.txt
```



FIG 12: INSTALL SATELLITE

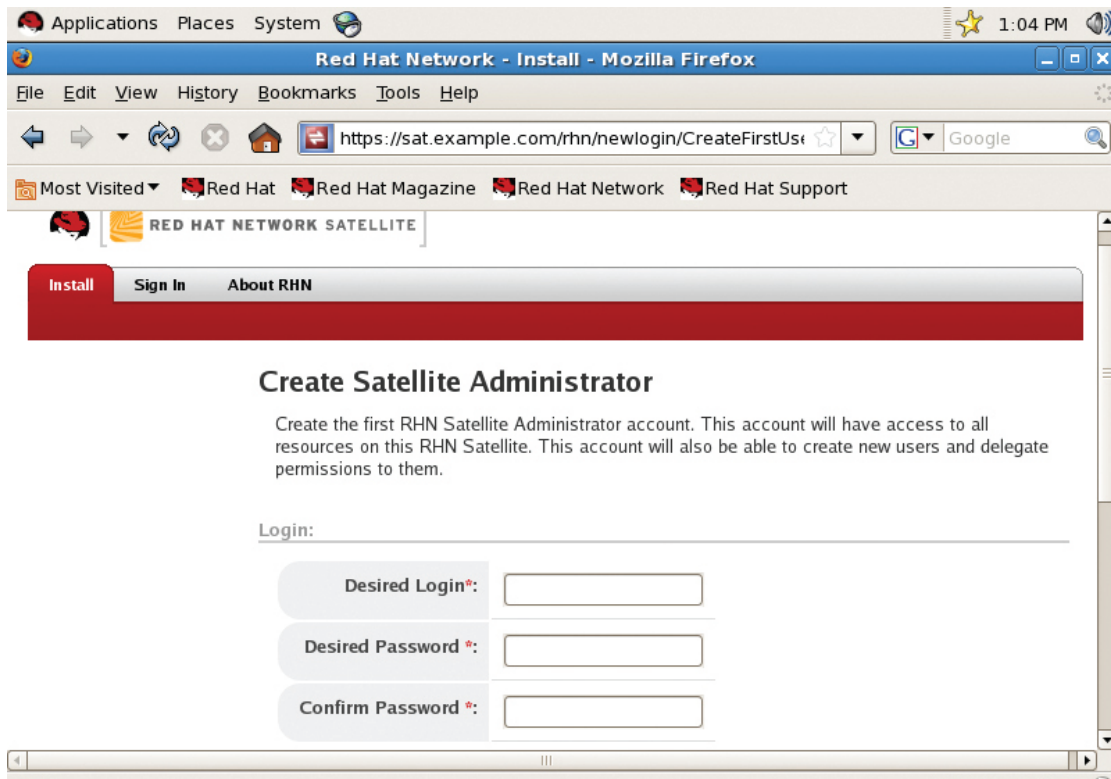


```
root@node1:~  
File Edit View Terminal Tabs Help  
[root@node1 ~]# ./install.pl --disconnected --answer-file=/tmp/answers.txt
```

3. Once installation has completed, open browser and navigate to <https://sat.example.com>.
4. Create Organization Administration account.
5. Configure Satellite settings.
6. Follow steps in Verifying Satellite Operation section.



FIG 13: SATELLITE INITIAL LOGIN



Back up Satellite database/configuration files

1. Disable rhn-database service.

```
# service rhn-database stop
```

2. Create backup directory.

```
# mkdir -p /tmp/sat-backup/db-backup; chown oracle /tmp/sat-backup/db-backup
```



FIG 14: SATELLITE DB BACKUP

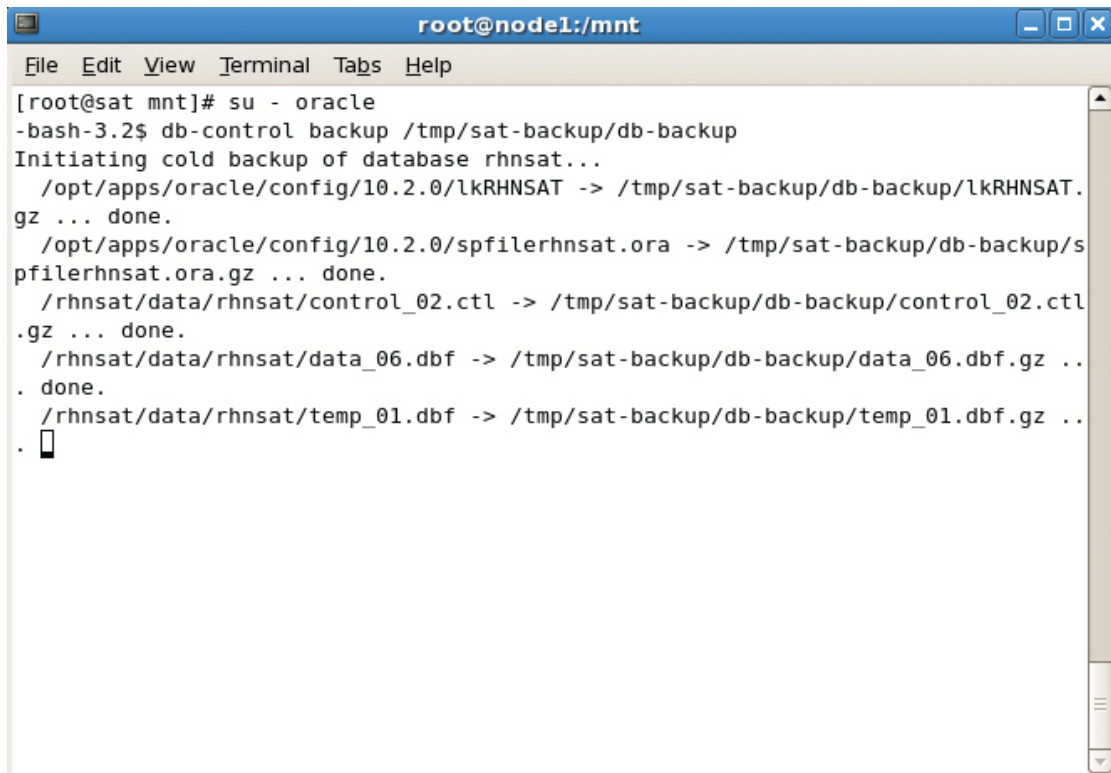
```
root@node1:/mnt
File Edit View Terminal Tabs Help
[root@sat mnt]# service rhn-database stop
Shutting down rhn-database: [ OK ]
[root@sat mnt]# mkdir -p /tmp/sat-backup/db-backup
[root@sat mnt]# chown oracle /tmp/sat-backup/db-backup
[root@sat mnt]#
```

3. Backup Satellite database.

```
# su - oracle
$ db-control backup /tmp/sat-backup/db-backup
$ exit
# service rhn-satellite stop
```



FIG 15A: SATELLITE DATABASE BACKUP

A terminal window titled 'root@node1:/mnt' showing the execution of a database backup script. The user switches to the 'oracle' user and runs 'db-control backup /tmp/sat-backup/db-backup'. The script initiates a cold backup of the 'rhnsat' database, compressing various files into a temporary backup directory. The files listed include configuration files, spfiles, control files, and datafiles.

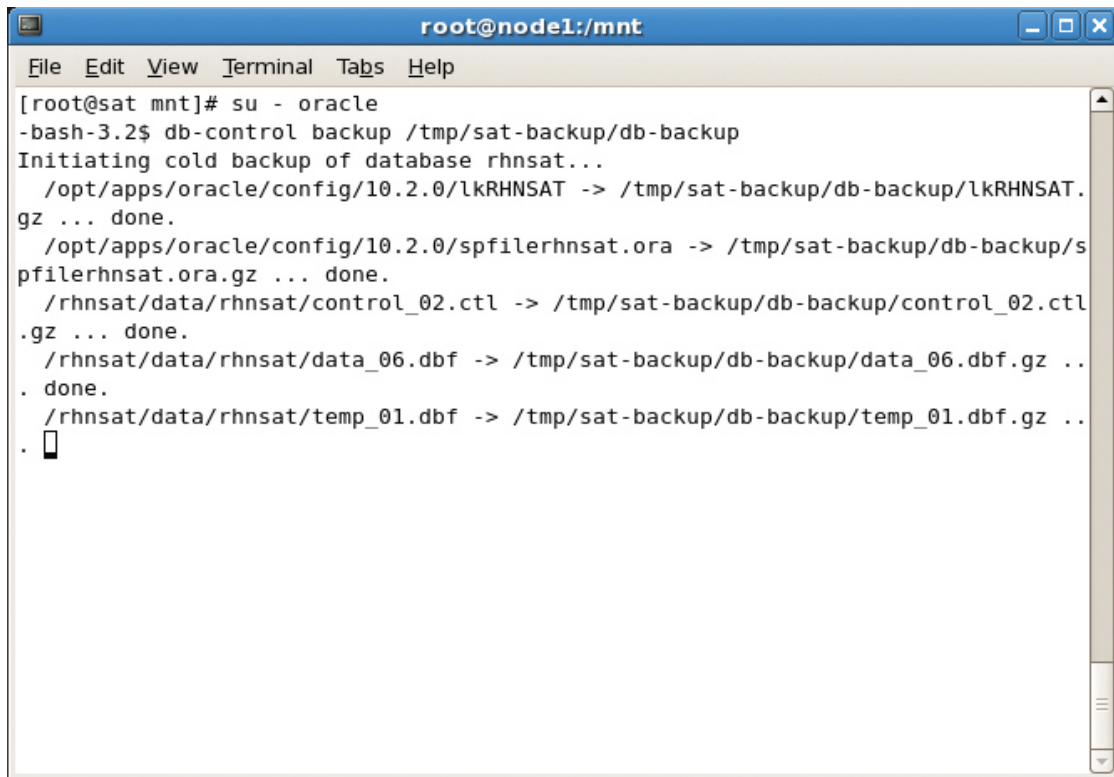
```
root@node1:/mnt
File Edit View Terminal Tabs Help
[root@sat mnt]# su - oracle
-bash-3.2$ db-control backup /tmp/sat-backup/db-backup
Initiating cold backup of database rhnsat...
/opt/apps/oracle/config/10.2.0/lkrhnsat -> /tmp/sat-backup/db-backup/lkrhnsat.
gz ... done.
/opt/apps/oracle/config/10.2.0/spfilerhnsat.ora -> /tmp/sat-backup/db-backup/s
pfilerhnsat.ora.gz ... done.
/rhnsat/data/rhnsat/control_02.ctl -> /tmp/sat-backup/db-backup/control_02.ctl
.gz ... done.
/rhnsat/data/rhnsat/data_06.dbf -> /tmp/sat-backup/db-backup/data_06.dbf.gz ..
. done.
/rhnsat/data/rhnsat/temp_01.dbf -> /tmp/sat-backup/db-backup/temp_01.dbf.gz ..
. █
```

#### 4. Back up critical satellite files.

```
# cp -a /var/www/html/pub /tmp/sat-backup/
# cp -a /root/ssl-build /tmp/sat-backup/
# cp /etc/jabberd/server.pem /tmp/sat-backup/
# cp -a /etc/httpd/conf/ssl.* /tmp/sat-backup/
```



FIG 16: CONFIG FILE BACKUP



```
root@node1:/mnt
File Edit View Terminal Tabs Help
[root@sat mnt]# su - oracle
-bash-3.2$ db-control backup /tmp/sat-backup/db-backup
Initiating cold backup of database rhnsat...
/opt/apps/oracle/config/10.2.0/lkrhnsat -> /tmp/sat-backup/db-backup/lkrhnsat.
gz ... done.
/opt/apps/oracle/config/10.2.0/spfilerhnsat.ora -> /tmp/sat-backup/db-backup/s
pfilerhnsat.ora.gz ... done.
/rhnsat/data/rhnsat/control_02.ctl -> /tmp/sat-backup/db-backup/control_02.ctl
.gz ... done.
/rhnsat/data/rhnsat/data_06.dbf -> /tmp/sat-backup/db-backup/data_06.dbf.gz ..
. done.
/rhnsat/data/rhnsat/temp_01.dbf -> /tmp/sat-backup/db-backup/temp_01.dbf.gz ..
. █
```

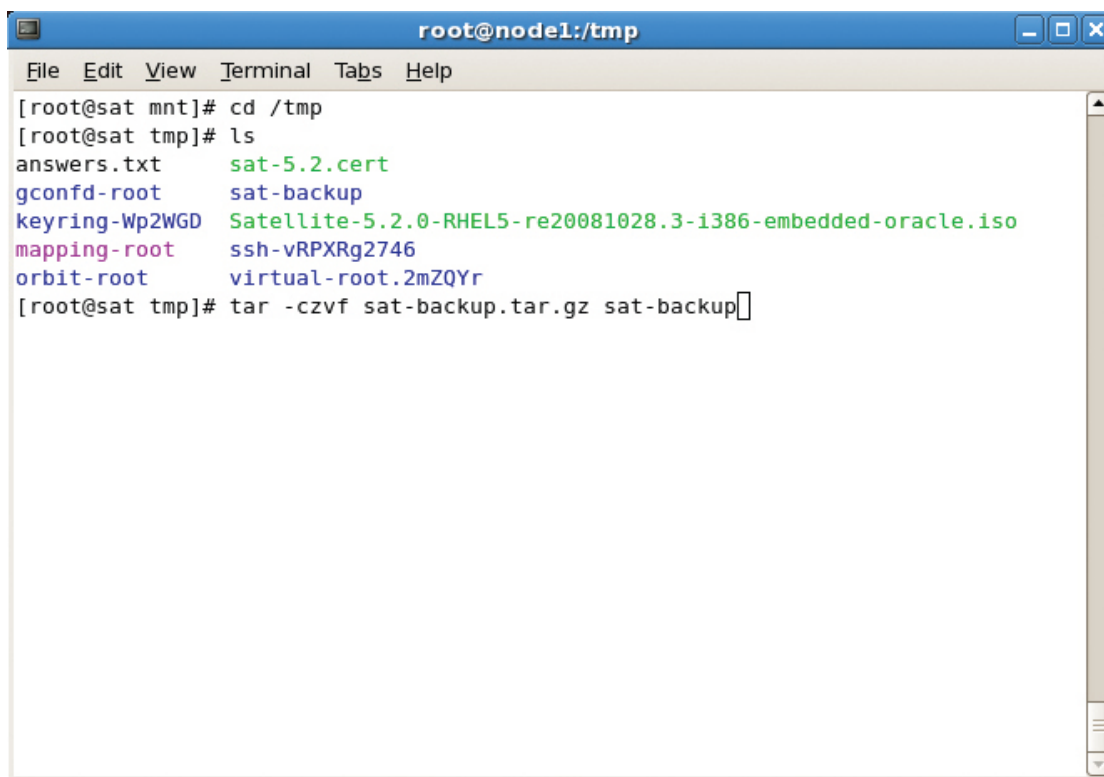
5. Archive and export backup files to Node 2.

```
# cd /tmp
# tar -czvf sat-backup.tar.gz sat-backup
# scp sat-backup.tar.gz node2.example.com:/tmp
```





FIG 17: ARCHIVE CONFIG FILES

A terminal window titled "root@node1:/tmp" with a menu bar (File, Edit, View, Terminal, Tabs, Help) and window control buttons. The terminal shows the following commands and output:

```
[root@sat mnt]# cd /tmp
[root@sat tmp]# ls
answers.txt      sat-5.2.cert
gconfd-root     sat-backup
keyring-Wp2WGD  Satellite-5.2.0-RHEL5-re20081028.3-i386-embedded-oracle.iso
mapping-root    ssh-vRPXRg2746
orbit-root      virtual-root.2mZQYr
[root@sat tmp]# tar -czvf sat-backup.tar.gz sat-backup
```



FIG 18: MIGRATE CONFIG FILES

```
root@node1:/tmp
File Edit View Terminal Tabs Help
[root@sat mnt]# cd /tmp
[root@sat tmp]# ls
answers.txt      sat-5.2.cert
gconfd-root     sat-backup
keyring-Wp2WGD  Satellite-5.2.0-RHEL5-re20081028.3-1386-embedded-oracle.iso
mapping-root    ssh-vRPXRg2746
orbit-root      virtual-root.2mZQYr
[root@sat tmp]# tar -czvf sat-backup.tar.gz sat-backup
```

### Disable Automatic Satellite Startup

Satellite Service management will be handled by Red Hat Cluster Suite. Therefore, it is necessary to disable Satellite Service from automatically starting on cluster nodes.

To achieve this, perform the following steps:

1. Turn off Satellite Service.

```
# service rhn-satellite stop
```

2. Disable Satellite and associated service.

```
# chkconfig rhn-satellite off
# chkconfig jabberd off; chkconfig rhn-database off; chkconfig osa-dispatcher off; chkconfig
taskomatic off; chkconfig tomcat5 off; chkconfig satellite-httpd off; chkconfig rhn-search
off
```



FIG 19: DISABLE SATELLITE AUTOMATIC STARTUP

A terminal window titled 'root@node1:/tmp' with a menu bar (File, Edit, View, Terminal, Tabs, Help). The terminal shows the following commands and output:

```
[root@sat mnt]# cd /tmp
[root@sat tmp]# ls
answers.txt      sat-5.2.cert
gconfd-root     sat-backup
keyring-Wp2WGD  Satellite-5.2.0-RHEL5-re20081028.3-i386-embedded-oracle.iso
mapping-root    ssh-vRPXRg2746
orbit-root      virtual-root.2mZQYr
[root@sat tmp]# tar -czvf sat-backup.tar.gz sat-backup
```

If monitoring functionality is employed in Satellite, it is disabled by using the following procedure:

1. Find and disable Monitoring and MonitoringScout.

```
# find /etc/rc.d/rc*.d -name '*Monitor*' -exec unlink {} \;
```

Confirm /rhnsat and /var/satellite mount points are not automatically mounted, by verifying contents of /etc/fstab.

It is also necessary to release shared resources from node 1.

1. Release shared IP address.

```
# ip addr del 192.168.1.30/24 dev bond0
```



FIG 20: DELETE VIRTUAL IP

```
root@localhost:/etc/sysconfig/network-scripts
File Edit View Terminal Tabs Help
[root@node1 network-scripts]# ip addr add 192.168.1.30/24 dev bond0
[root@node1 network-scripts]# ip addr list dev bond0
4: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether 00:17:de:36:26:01 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.10/24 brd 192.168.1.255 scope global bond0
    inet 192.168.1.30/24 scope global secondary bond0
    inet6 fe80::217:deff:fe36:2601/64 scope link
        valid_lft forever preferred_lft forever
[root@node1 network-scripts]# ip addr del 192.168.1.30/24 dev bond0
[root@node1 network-scripts]# ip addr list dev bond0
4: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether 00:17:de:36:26:01 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.10/24 brd 192.168.1.255 scope global bond0
    inet6 fe80::217:deff:fe36:2601/64 scope link
        valid_lft forever preferred_lft forever
[root@node1 network-scripts]#
```

2. Unmount shared storage.

```
# umount /rhnsat
# umount /var/satellite
```

### Satellite installation (Node 2)

Once Satellite service on Node 1 has been verified to function properly, meaning clients can be subscribed to satellite, receive updates, etc., it is necessary to export critical database and configuration files exported and verify automatic startup has been disabled; then, installation of Satellite on Node 2 can begin.

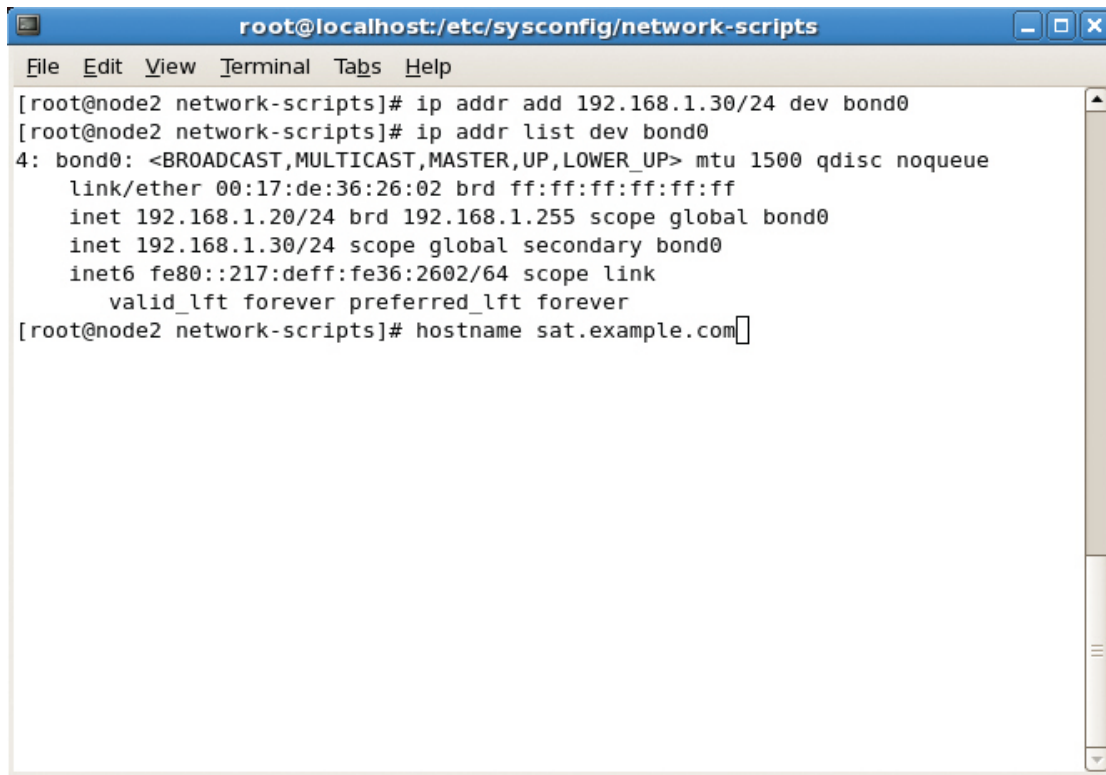
For Satellite installation on Node 2, perform the following:

1. Configure virtual IP address.

```
# ip addr add 192.168.1.30 dev bond0
```



FIG 21: CREATE VIRTUAL IP



```
root@localhost:/etc/sysconfig/network-scripts
File Edit View Terminal Tabs Help
[root@node2 network-scripts]# ip addr add 192.168.1.30/24 dev bond0
[root@node2 network-scripts]# ip addr list dev bond0
4: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether 00:17:de:36:26:02 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.20/24 brd 192.168.1.255 scope global bond0
    inet 192.168.1.30/24 scope global secondary bond0
    inet6 fe80::217:deff:fe36:2602/64 scope link
        valid_lft forever preferred_lft forever
[root@node2 network-scripts]# hostname sat.example.com
```

2. Insert and Mount Satellite Installation Media.

```
# mount -o loop /tmp/Satellite-5.2-RHEL5.iso /mnt
```

3. Begin installation.

```
# ./install.pl -disconnected -answer-file=/tmp/answers.txt
```

Note: Do not navigate to <https://sat.example.com> to create admin login information. This information will be imported from Satellite installation on Node 1.

4. Stop Satellite Service.

```
# service rhn-satellite stop
```

5. Install Satellite SSL Certificate.

```
# rpm -Uvh http://node1.example.com/pub/rhn-org-trusted-ssl-cert-1.0-1.noarch.rpm
```



FIG 22: INSTALL SATELLITE CERTIFICATE

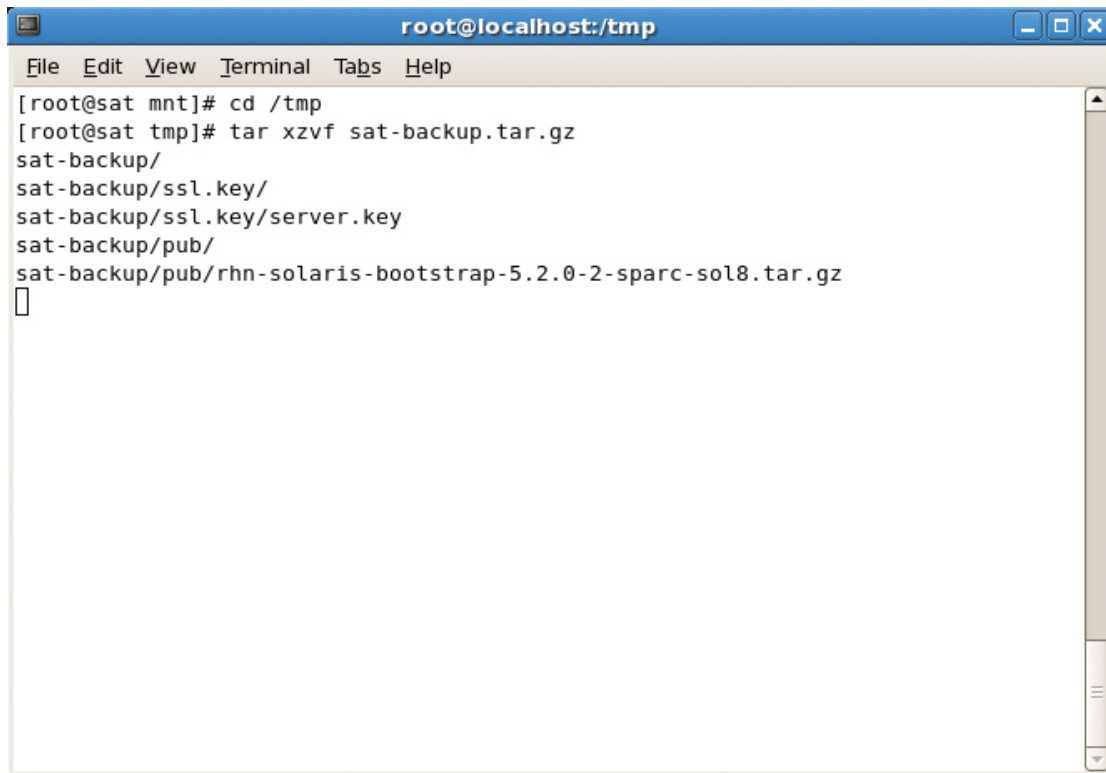
```
root@localhost:/mnt
File Edit View Terminal Tabs Help
[root@sat mnt]# service rhn-satellite stop
Shutting down rhn-satellite...
Stopping rhn-search...
Stopped rhn-search.
Stopping satellite-httpd:          [ OK ]
Stopping tomcat5:                  [ OK ]
Stopping RHN Taskomatic...
Stopped RHN Taskomatic.
Shutting down osa-dispatcher:      [ OK ]
Shutting down rhn-database:        [ OK ]
Shutting down Jabber router:       [ OK ]
Done.
[root@sat mnt]# rpm -ivh http://node1.example.com/pub/rhn-org-trusted-ssl-cert-1.0-1.noarch.rpm
```

6. Extract Satellite configuration files and install Satellite configuration files.

```
# cd /tmp; tar xzvf sat-backup.tar.gz
# cd sat-backup
```



FIG 23: EXTRACT CONFIG FILES

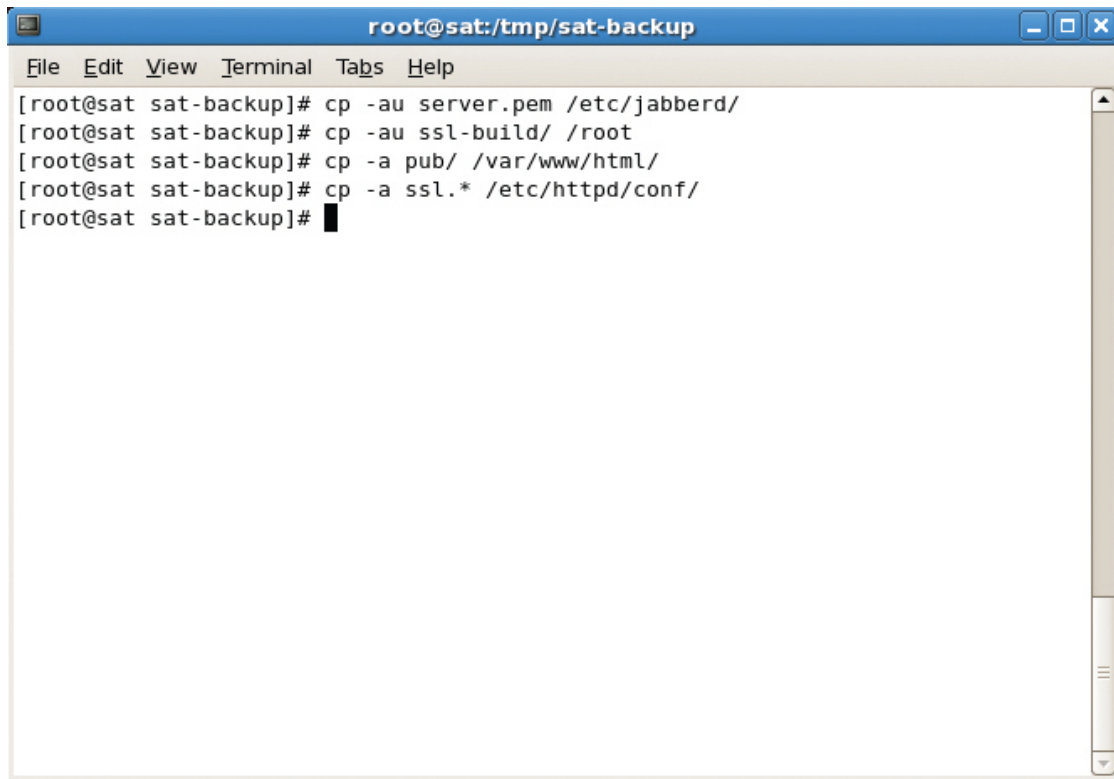
A terminal window titled "root@localhost:/tmp" with a menu bar (File, Edit, View, Terminal, Tabs, Help). The terminal shows the following commands and output:

```
[root@sat mnt]# cd /tmp
[root@sat tmp]# tar xzvf sat-backup.tar.gz
sat-backup/
sat-backup/ssl.key/
sat-backup/ssl.key/server.key
sat-backup/pub/
sat-backup/pub/rhn-solaris-bootstrap-5.2.0-2-sparc-sol8.tar.gz
█
```

```
# cp server.pem /etc/jabberd
# cp -au ssl-build/ /root
# cp -a pub/ /var/www/html/
# cp -a etc/httpd/conf/* /etc/httpd/conf/
```



FIG 24: INSTALL SATELLITE CONFIG FILES



```
root@sat:/tmp/sat-backup
File Edit View Terminal Tabs Help
[root@sat sat-backup]# cp -au server.pem /etc/jabberd/
[root@sat sat-backup]# cp -au ssl-build/ /root
[root@sat sat-backup]# cp -a pub/ /var/www/html/
[root@sat sat-backup]# cp -a ssl.* /etc/httpd/conf/
[root@sat sat-backup]#
```

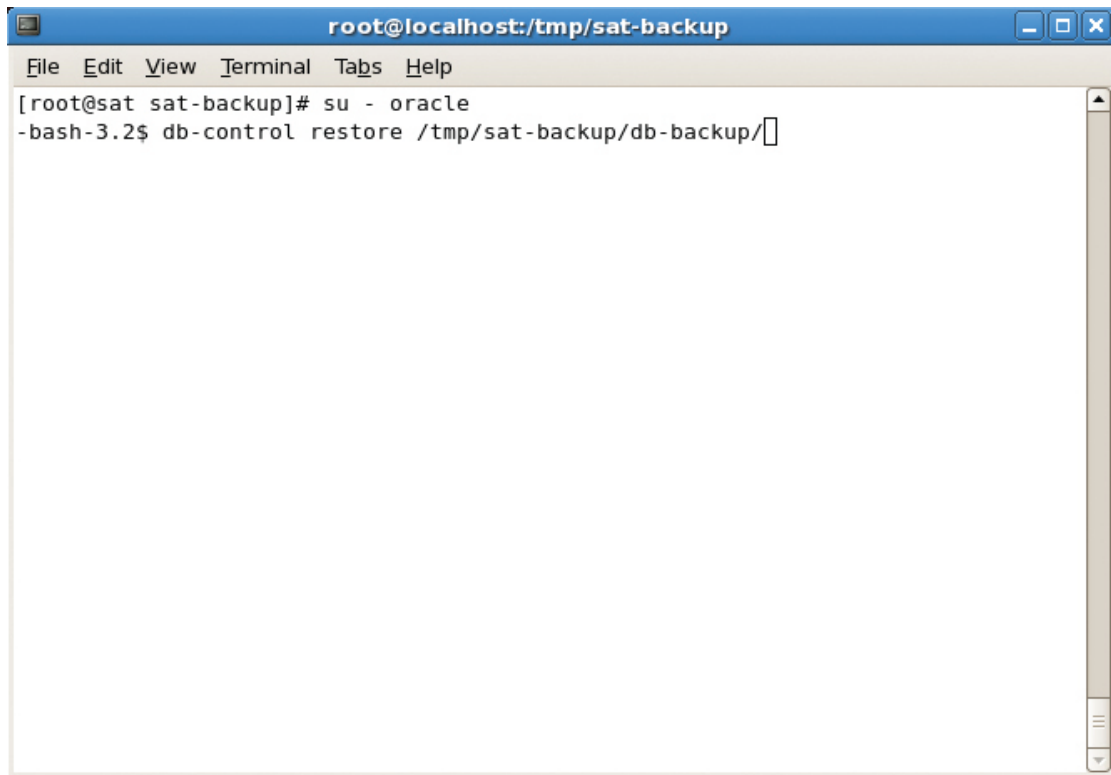
### 7. Restore Satellite database.

```
# su - oracle
$ db-control restore /tmp/sat-backup/db-backup/
$ exit
```





FIG 25A: RESTORE DATABASE

A terminal window titled 'root@localhost:/tmp/sat-backup' with a menu bar (File, Edit, View, Terminal, Tabs, Help). The terminal shows the user switching to the 'oracle' user and running the command 'db-control restore /tmp/sat-backup/db-backup/'.

```
root@localhost:/tmp/sat-backup
File Edit View Terminal Tabs Help
[root@sat sat-backup]# su - oracle
-bash-3.2$ db-control restore /tmp/sat-backup/db-backup/
```

### Verify Satellite operates properly

1. Start Satellite Service.

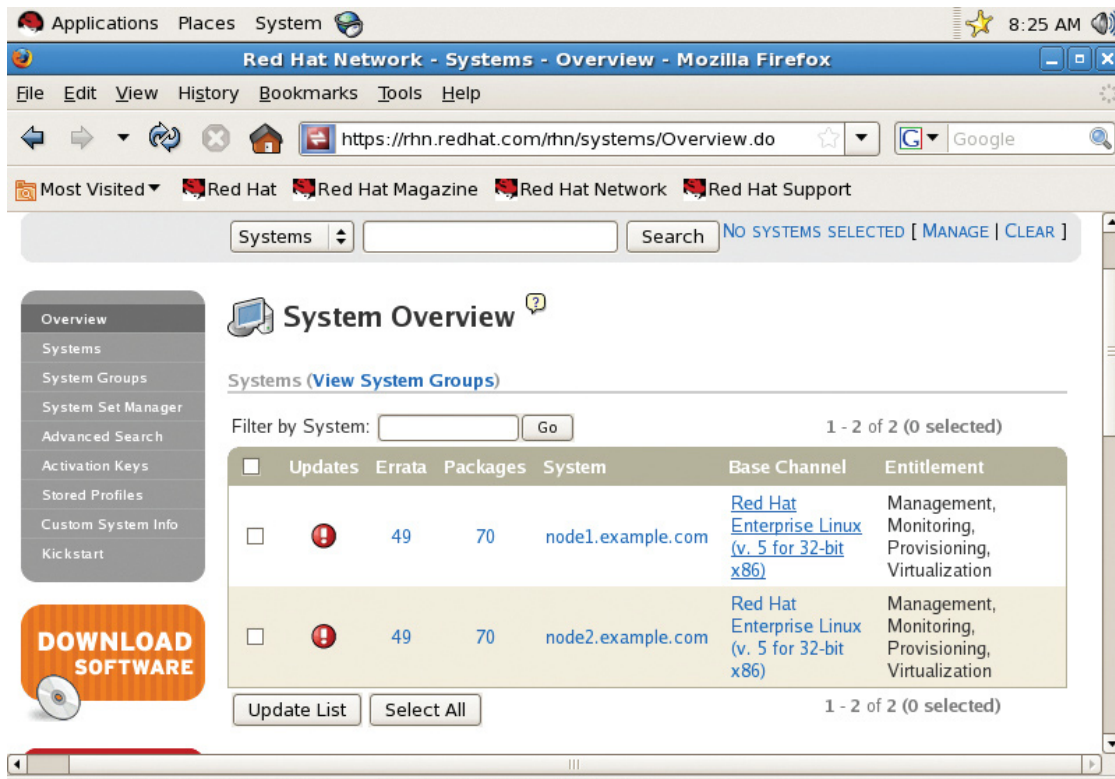
```
# service rhn-satellite start
```

2. Open browser, navigate to <https://sat.example.com>. Login with credentials created during node1 Satellite installation procedure.
3. Verify clients previously registered are present in Satellite.
4. Navigate to Satellite Tools Section in gray horizontal navigation bar.
5. Select Satellite Configuration in gray operation box on left.
6. Click monitoring link in main section of screen below RHN Satellite Configuration paragraph.
7. Enable monitoring on Satellite Server.
8. Close web browser.
9. Restart Satellite Server.

```
# service rhn-satellite restart
```



FIG 26A: ADD CLUSTERING CHANNEL SUBSCRIPTION



## RED HAT CLUSTER SUITE

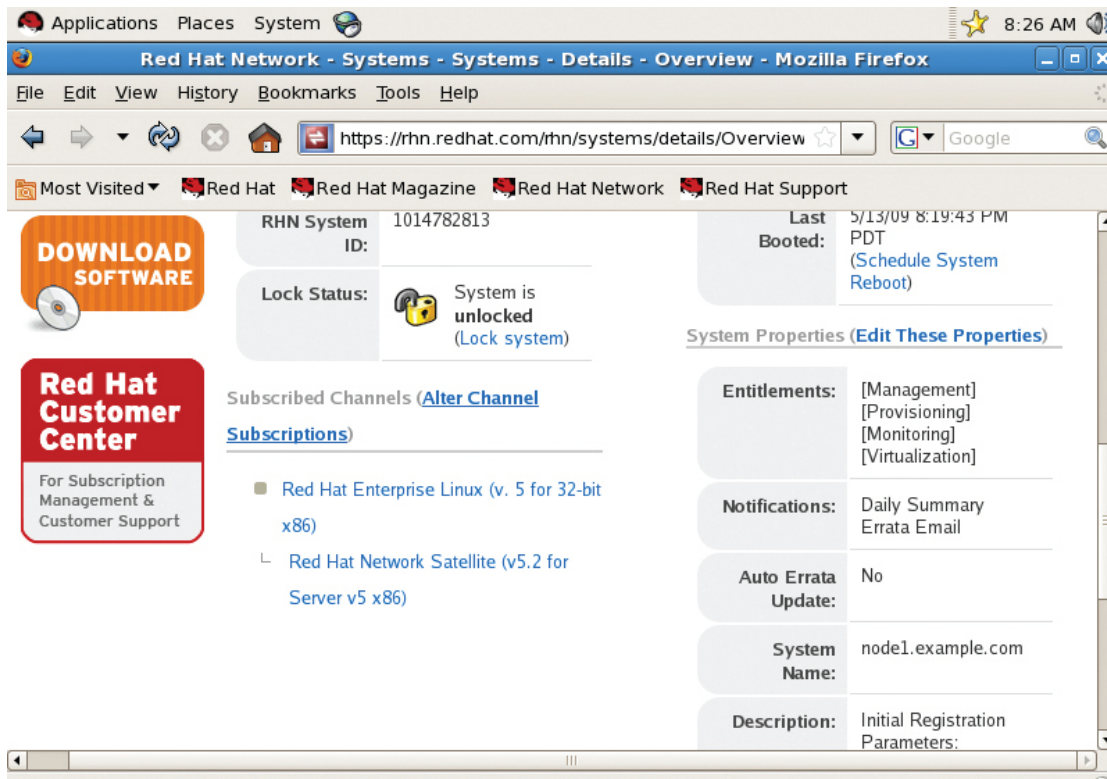
### Red Hat Cluster Suite Subscription

Red Hat Cluster suite was installed by first subscribing cluster nodes to Red Hat Enterprise Linux Clustering (v. 5 for 64-bit x86\_64).

1. From the Red Hat Network website <http://rhn.redhat.com> click Systems in the red navigation bar.
2. Select system(s) that will be participants in clustering



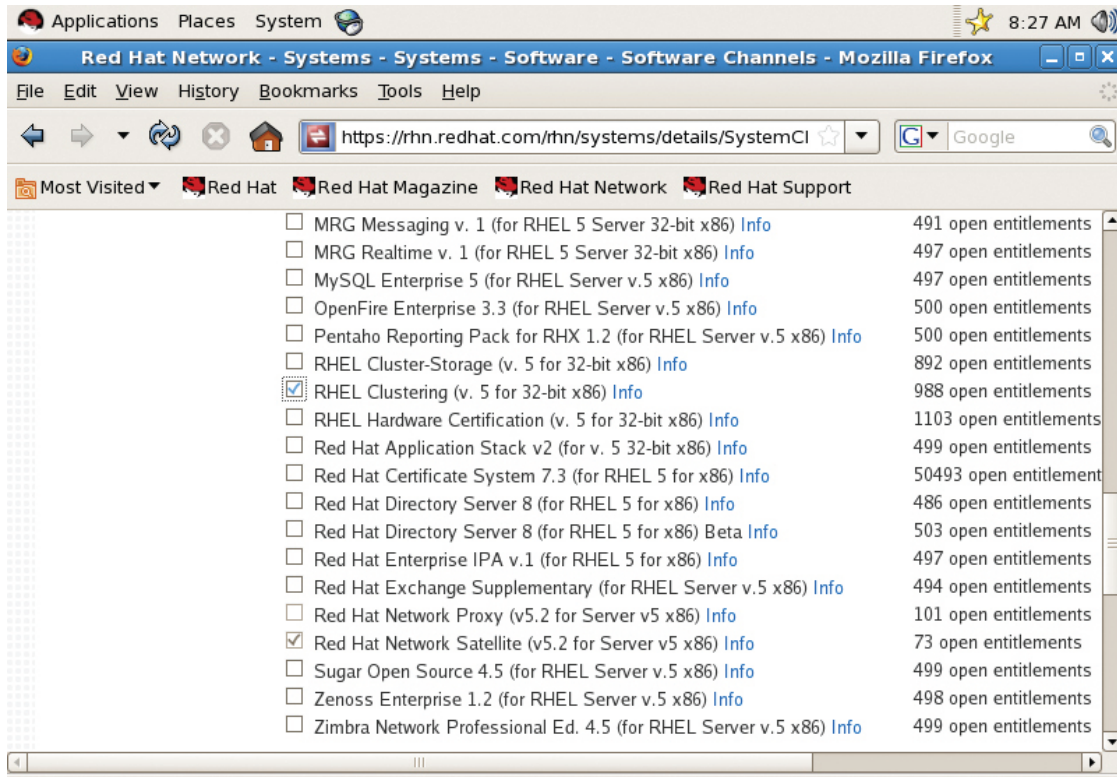
FIG 27B: ADD CLUSTERING CHANNEL SUBSCRIPTION



3. Select Alter Channel Subscription in Subscribed Channels area.



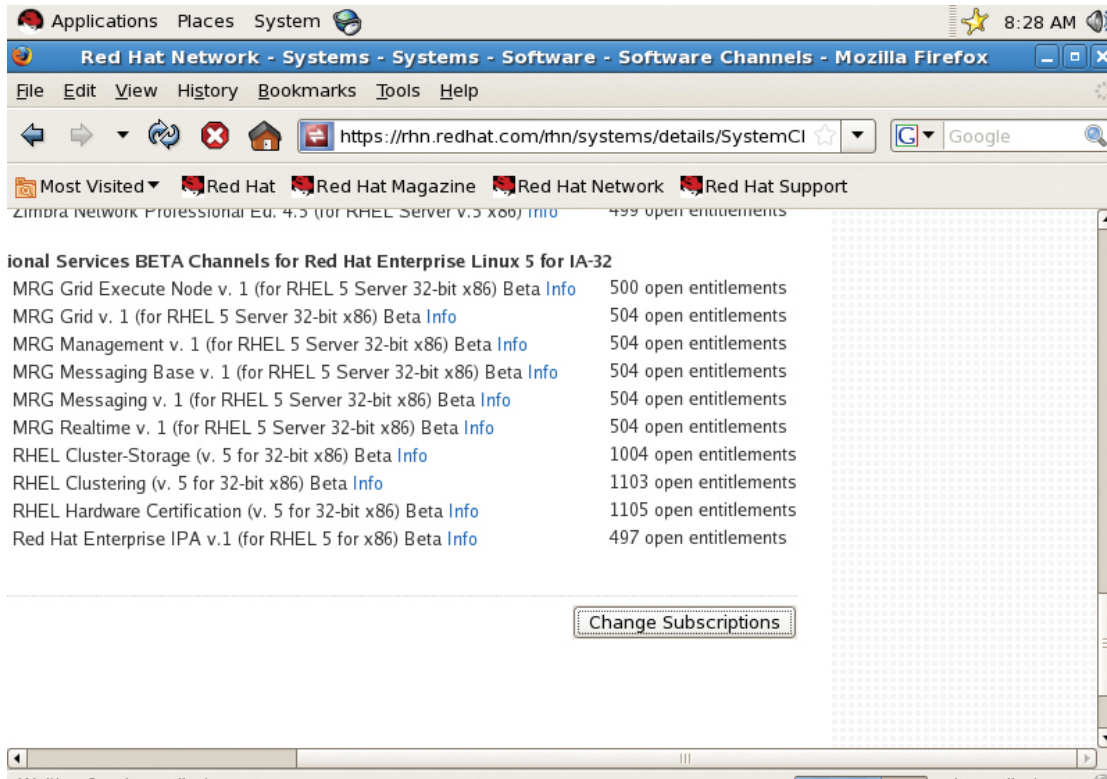
FIG 28C: SELECT RHEL CLUSTERING



4. Check RHEL Clustering (v. 5 for 64-bit x86\_64) under Additional Services Channels.



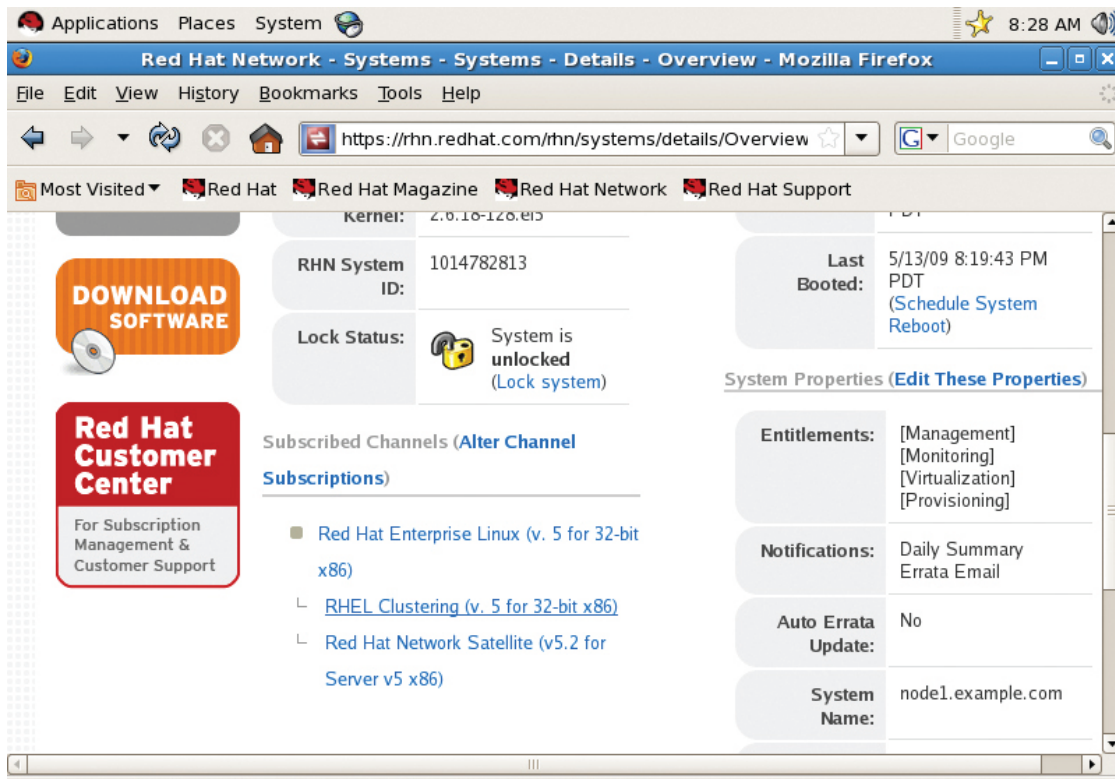
FIG 29D: CONFIRM CHANNEL SUBSCRIPTION



5. Scroll to bottom and select Confirm.



FIG 30E: VERIFY NEW SUBSCRIPTION



6. Click Details and verify system is subscribed to channel under Subscribed Channels area.





FIG 31: INSTALL CLUSTERING PACKAGES

```

root@sat:~
File Edit View Terminal Tabs Help
=====
Package      Arch   Version      Repository      Size
=====
Installing:
Cluster_Administration-en-US
  Cluster_Administration-en-US  noarch  5.2-1        rhel-i386-server-cluster-5  1.7 M
rgmanager     i386   2.0.46-1.el5_3.3  rhel-i386-server-cluster-5  287 k
system-config-cluster
  system-config-cluster         noarch  1.0.55-1.0    rhel-i386-server-cluster-5  291 k
Installing for dependencies:
cman          i386   2.0.98-1.el5_3.1  rhel-i386-server-5          609 k
openais      i386   0.80.3-22.el5_3.4  rhel-i386-server-5          375 k
perl-Net-Telnet  noarch  3.03-5           rhel-i386-server-5          56 k
pexpect      noarch  2.3-1.el5        rhel-i386-server-5          218 k

Transaction Summary
=====
Install      7 Package(s)
Update       0 Package(s)
Remove       0 Package(s)

Total download size: 3.5 M
Is this ok [y/N]: █

```

Firewall settings were disabled, along with SELinux set to disabled. Kernel dump was disabled, and system was not immediately registered to Red Hat Network.

## RED HAT CLUSTER SUITE INSTALLATION

For this engagement, system-config-cluster application was used to install Cluster Suite on node systems. For additional methods of installation please refer to: [http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5.2/html/Cluster\\_Administration/index.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Cluster_Administration/index.html)

Red Hat Enterprise Linux on nodes was built according to the guidelines above for Red Hat Enterprise Linux installation.

1. Install Cluster Software suite.

```
# yum -y groupinstall clustering
```



FIG 32: FDISK QUORUM DISK

```
root@node2:~  
File Edit View Terminal Tabs Help  
[root@node2 ~]# fdisk -l /dev/sdb  
  
Disk /dev/sdb: 104 MB, 104857600 bytes  
4 heads, 50 sectors/track, 1024 cylinders  
Units = cylinders of 200 * 512 = 102400 bytes  
  
   Device Boot      Start         End      Blocks   Id  System  
/dev/sdb1            1          1024      102375    83  Linux  
[root@node2 ~]#
```

If installing Satellite in disconnected mode:

1. Insert and mount Red Hat Enterprise Linux 5 installation media.

```
# mount /dev/cdrom /mnt
```

2. Create clustering repository file.

```
# vi /etc/yum.repos.d/cluster.repo  
[redhat-cdrepo]  
name=Red Hat Enterprise Linux CD Repo  
baseurl=file:///mnt/Clustering  
enabled=1  
gpgcheck=0  
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
```





FIG 33: CREATE QUORUM DISK

```
root@node2:~  
File Edit View Terminal Tabs Help  
[root@node2 ~]# mkqdisk -c /dev/sdb1 -l quorum  
mkqdisk v0.6.0  
Writing new quorum disk label 'quorum' to /dev/sdb1.  
WARNING: About to destroy all data on /dev/sdb1; proceed [N/y] ? y  
Initializing status block for node 1...  
Initializing status block for node 2...  
Initializing status block for node 3...  
Initializing status block for node 4...  
Initializing status block for node 5...  
Initializing status block for node 6...  
Initializing status block for node 7...  
Initializing status block for node 8...  
Initializing status block for node 9...  
Initializing status block for node 10...  
Initializing status block for node 11...  
Initializing status block for node 12...  
Initializing status block for node 13...  
Initializing status block for node 14...  
Initializing status block for node 15...  
Initializing status block for node 16...  
[root@node2 ~]#
```

3. Clean yum cache and install clustering group.

```
# yum clean all  
# yum groupinstall clustering
```

4. The following packages not found in Clustering group are also required:

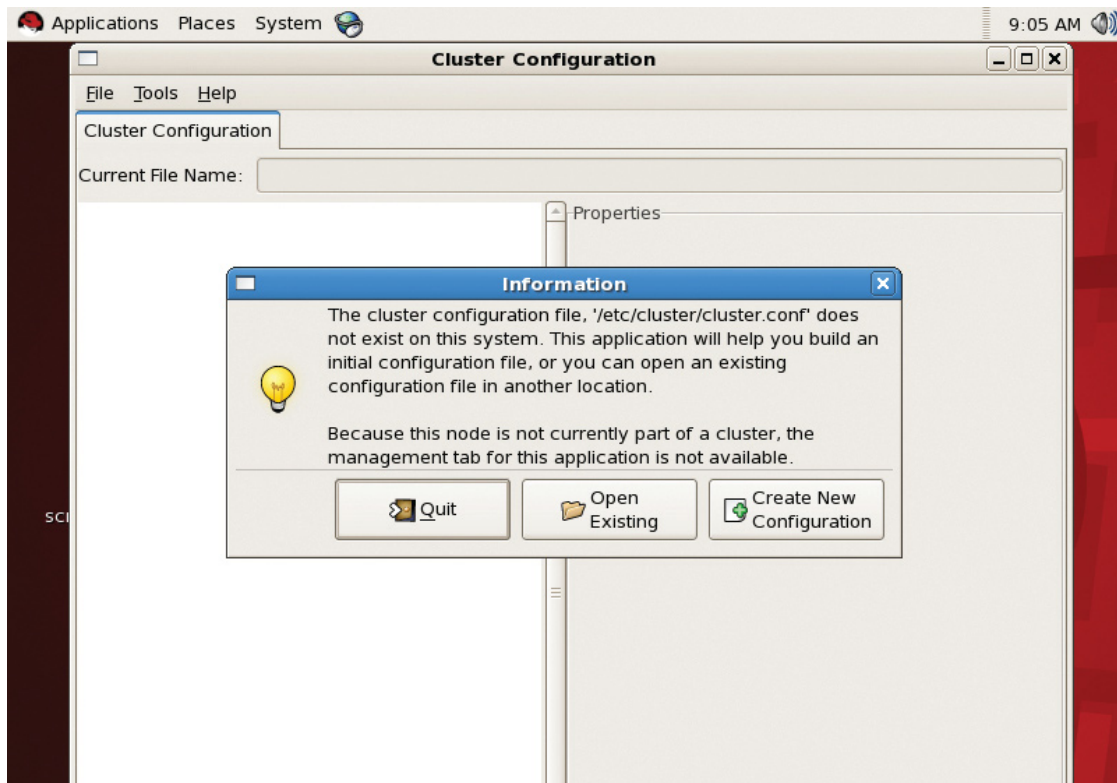
```
system-config-cluster  
rgmanager  
cman  
cluster-cim  
cluster-snmp
```

#### Configure quorum disk

1. Verify though fdisk quorum partition is available.  
# fdisk -l /dev/sdd1



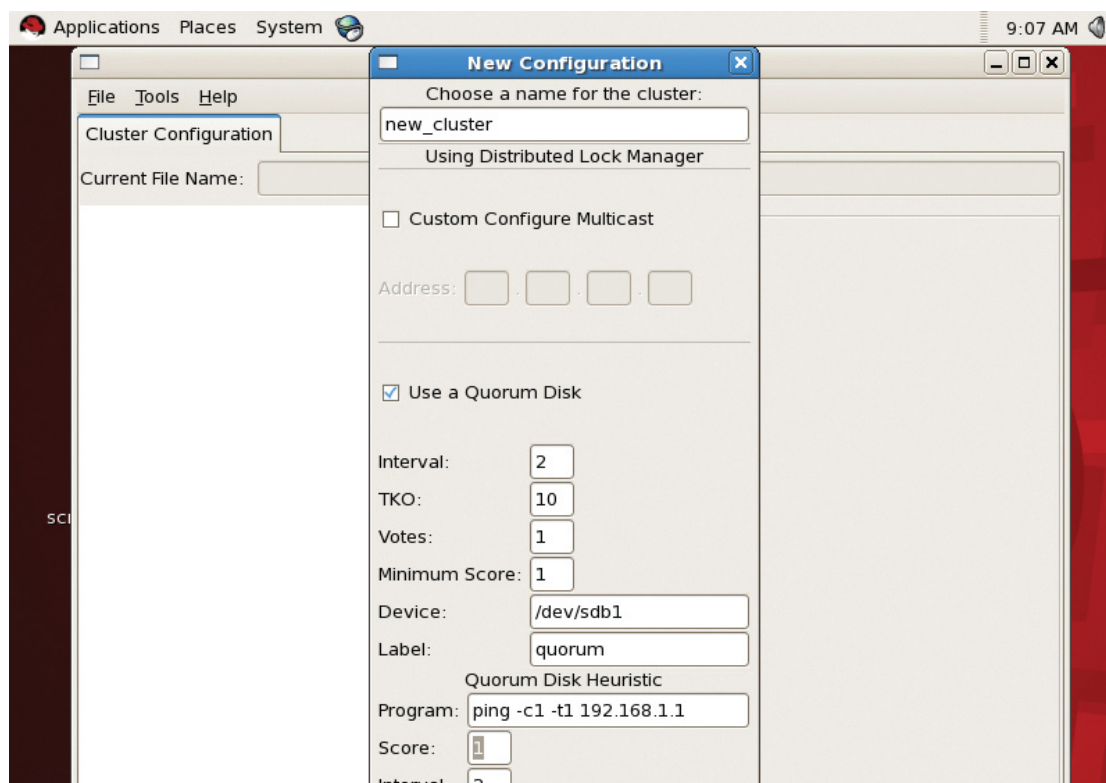
FIG 34: CREATE NEW CLUSTER



2. Create quorum disk.



FIG 35: INITIAL CLUSTER CONFIGURATION



### Configure cluster nodes

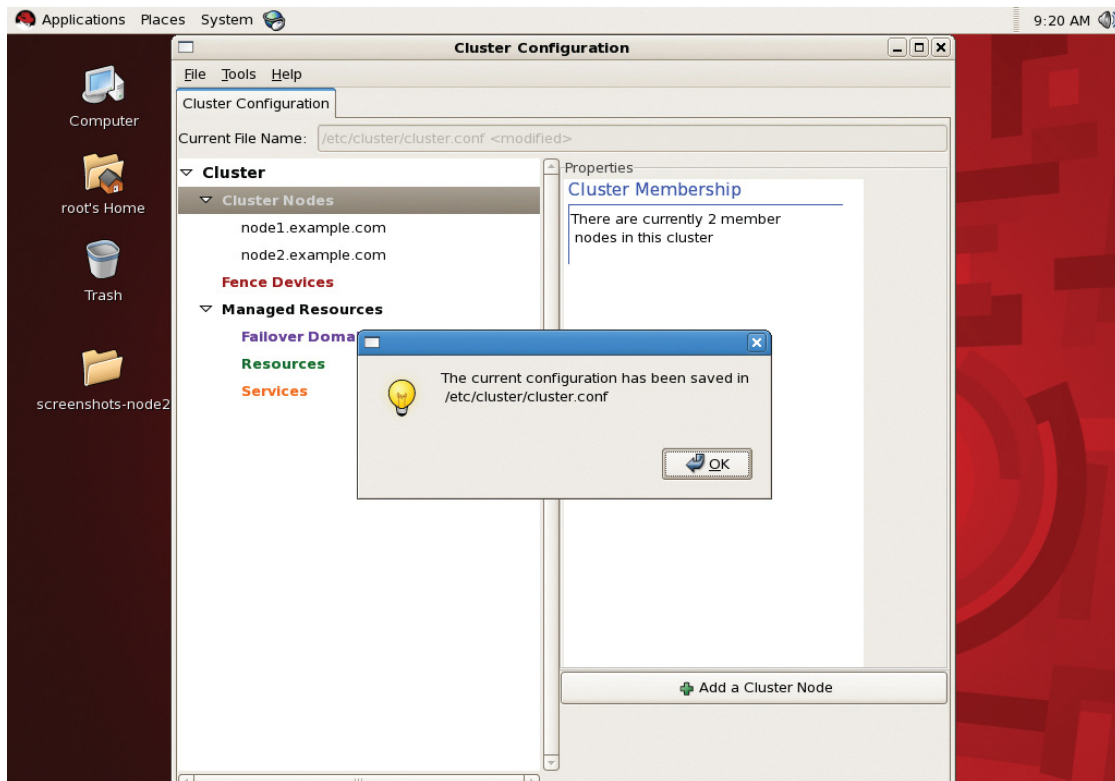
1. Start cluster configuration software.

```
# system-config-cluster
```

Select "Create New Configuration."



FIG 36: SAVE CLUSTER CONFIGURATION



2. Assign name to cluster and cluster details.

If using a quorum disk, check Use a quorum disk box. Insert device name and/or label of quorum disk. If you require specific factors to determine node health, enter the command to verify those factors in program under Quorum Disk Heuristics.



FIG 37: EDIT CLUSTER.CONFIG

```
root@node2:~
File Edit View Terminal Tabs Help
<?xml version="1.0" ?>
<cluster config_version="2" name="new_cluster">
  <quorumd device="/dev/sdd1" interval="2" label="quorum" min_score="1" tk
o="10" votes="1">
    <heuristic interval="2" program="ping -c1 -t1 192.168.1.1" score
="1"/>
  </quorumd>
  <fence_daemon post_fail_delay="0" post_join_delay="3"/>
  <clusternodes/>
  <cman expected_votes="3"/>
  <fencedevices/>
  <rm>
    <failoverdomains/>
    <resources/>
  </rm>
</cluster>
~
~
~
~
~
~
```

For a two node cluster, the following configurations were used for quorum disk:

```
Interval: 2
Votes: 1
TKO: 10
Minimum Score: 1
Device: /dev/sdb1
Label: quorum
Heuristics
Path to Program: ping -c1 -t1 192.168.1.1
Interval: 2
Score: 1
```

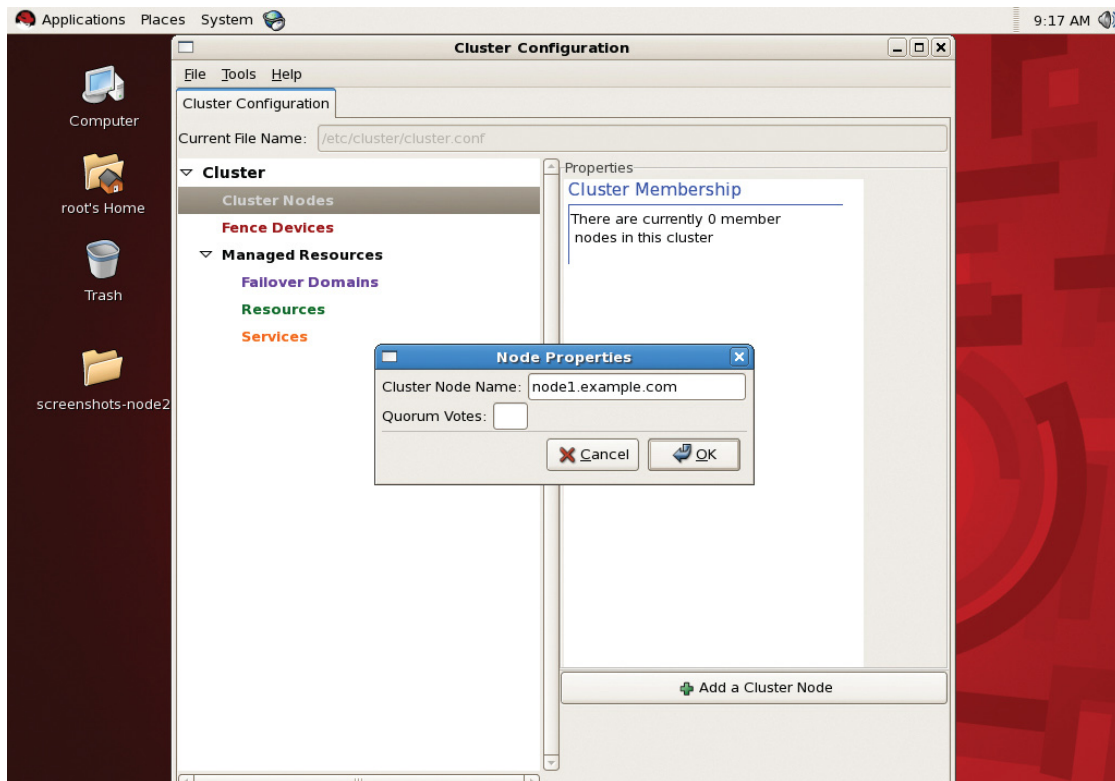
The above heuristic program is testing network connectivity of the nodes to the gateway IP address.

3. Click File --> Click Save.

4. Accept default file name of /etc/cluster/cluster.conf



FIG 38: ADD CLUSTER NODES



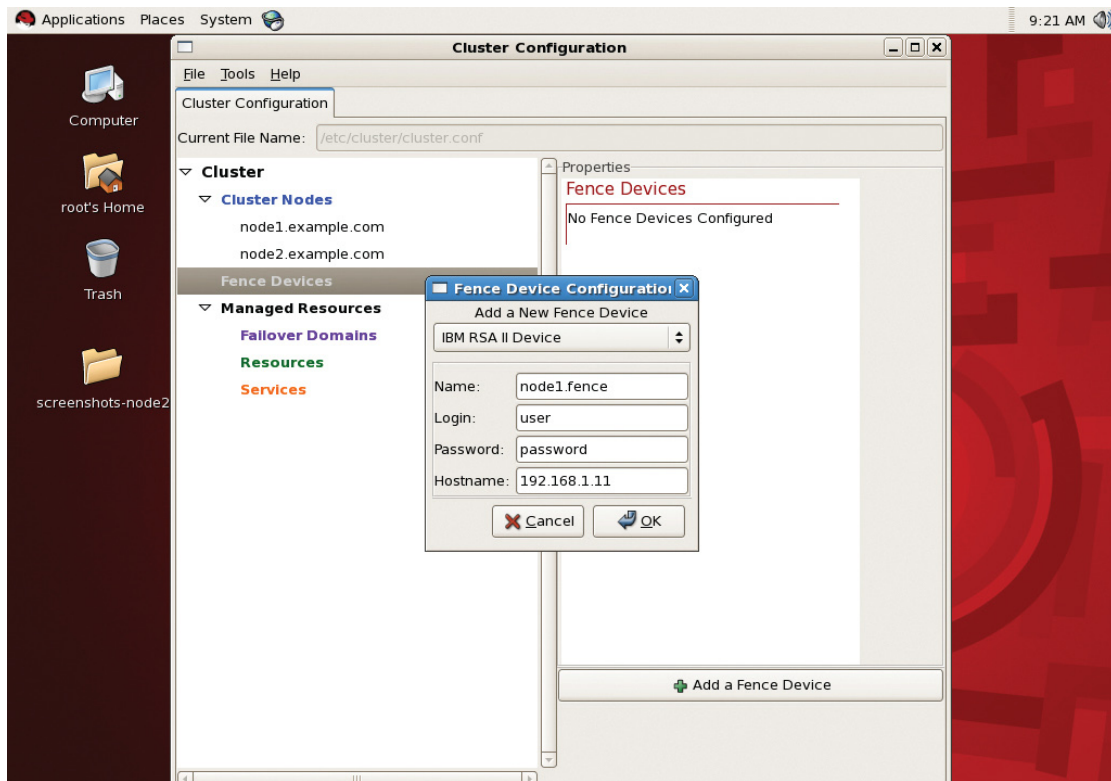
Additionally, manually add the following to <cman> section of /etc/cluster/cluster.conf :

```
expected_votes="3"
```

Increment config version number by one. Save and exit.



FIG 39: CREATE FENCE DEVICES

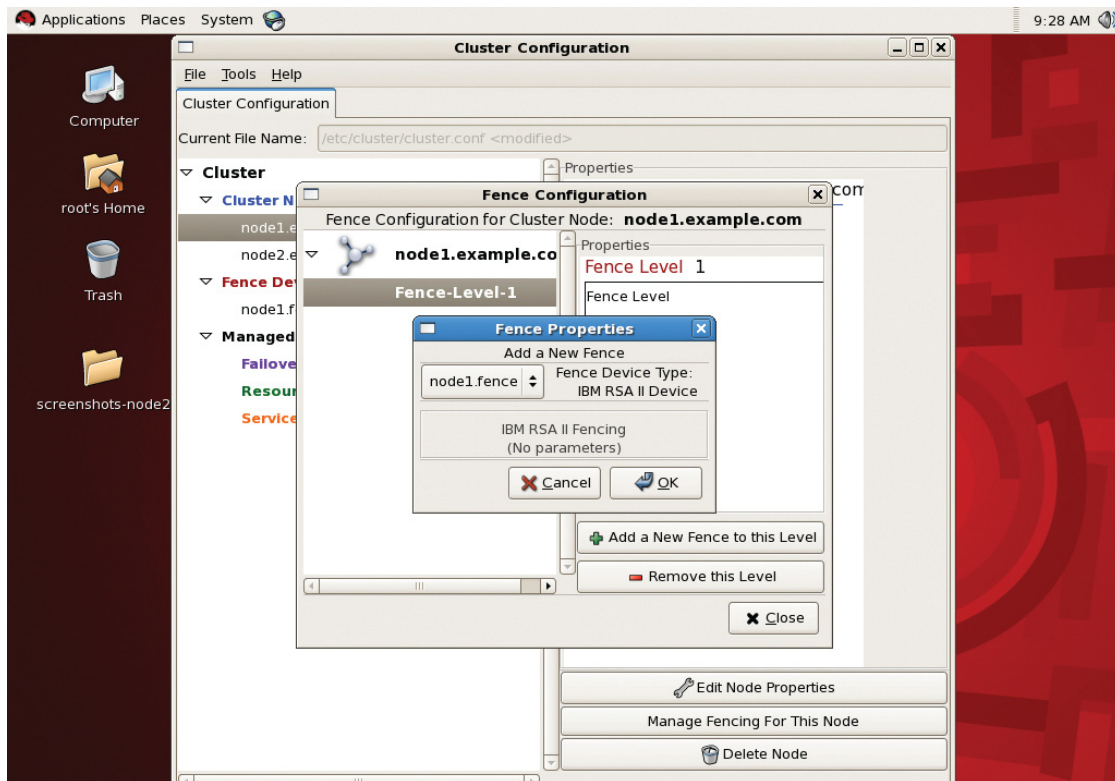


##### 5. Add Cluster Nodes.

- Click "Cluster Nodes" in left pane.
- Click "Add a Cluster Node" in right pane.
- Enter cluster name (FQDN hostname).
- quorum votes will be set automatically to 1.



FIG 40: ADD FENCE DEVICES TO NODES



6. Click Save.

Repeat above steps for each additional node in cluster.

### Add Fencing Device

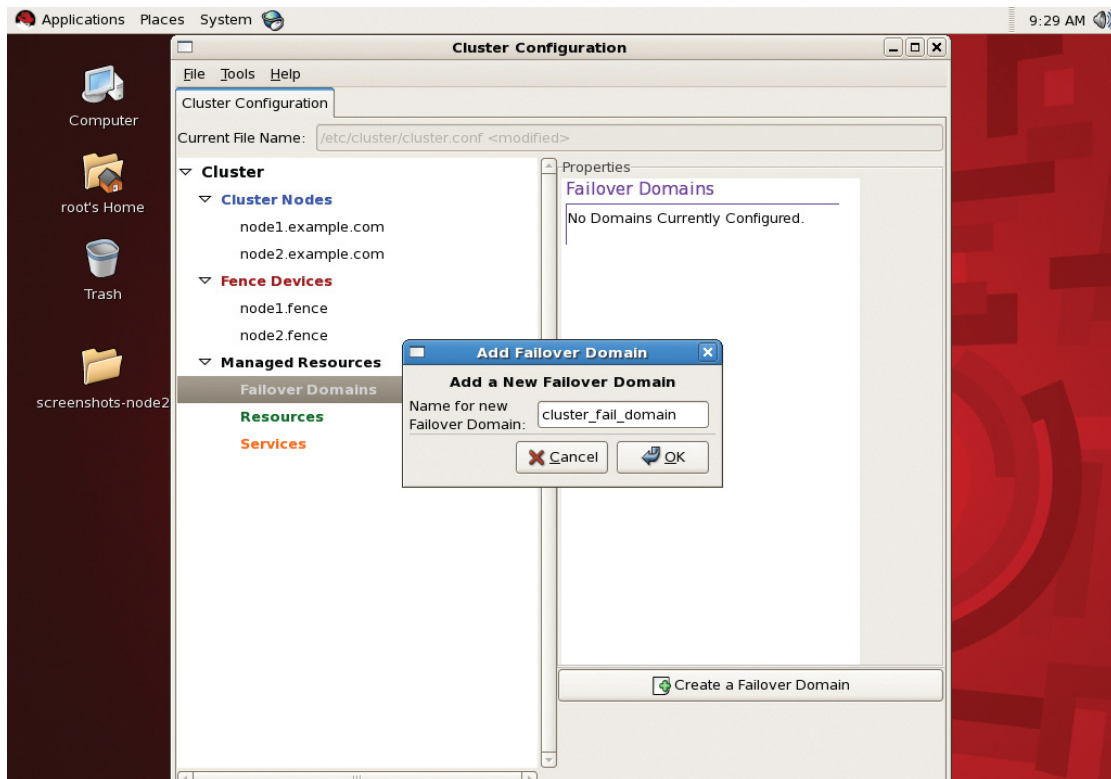
- Click "Fence Devices" in left pane.
- Click "Add a Fence Device" at bottom right of window.
- Select "IBM RSAll Device" in drop down menu.
- Fill required information (Name, login, password, hostname).
- Click OK.

Repeat the above procedure for additional nodes in cluster.





FIG 41: CREATE FAILOVER DOMAIN



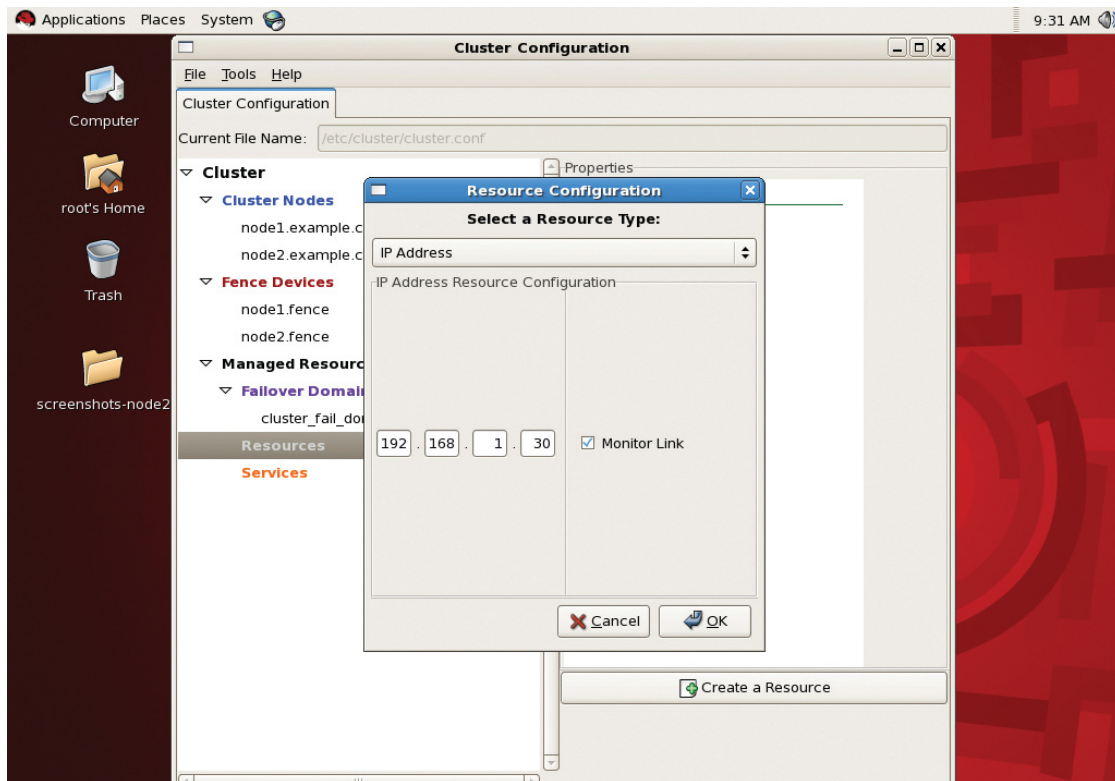
### Add Fence Devices to node

- Click first cluster node in left pane.
- In right pane, click "Manage Fencing for this node."
- In left pane of new window, click "Add a new Fence Level."
- In right pane, click "Fence-Level-1."
- In left pane, click "Add new fence to this level."
- Select newly created fence in drop down menu.
- Click Close.
- Click File --> Save.

Repeat above procedure for each additional node in cluster



FIG 42: CREATE IP RESOURCE



## CONFIGURE CLUSTER PROPERTIES

### Configuring failover domains

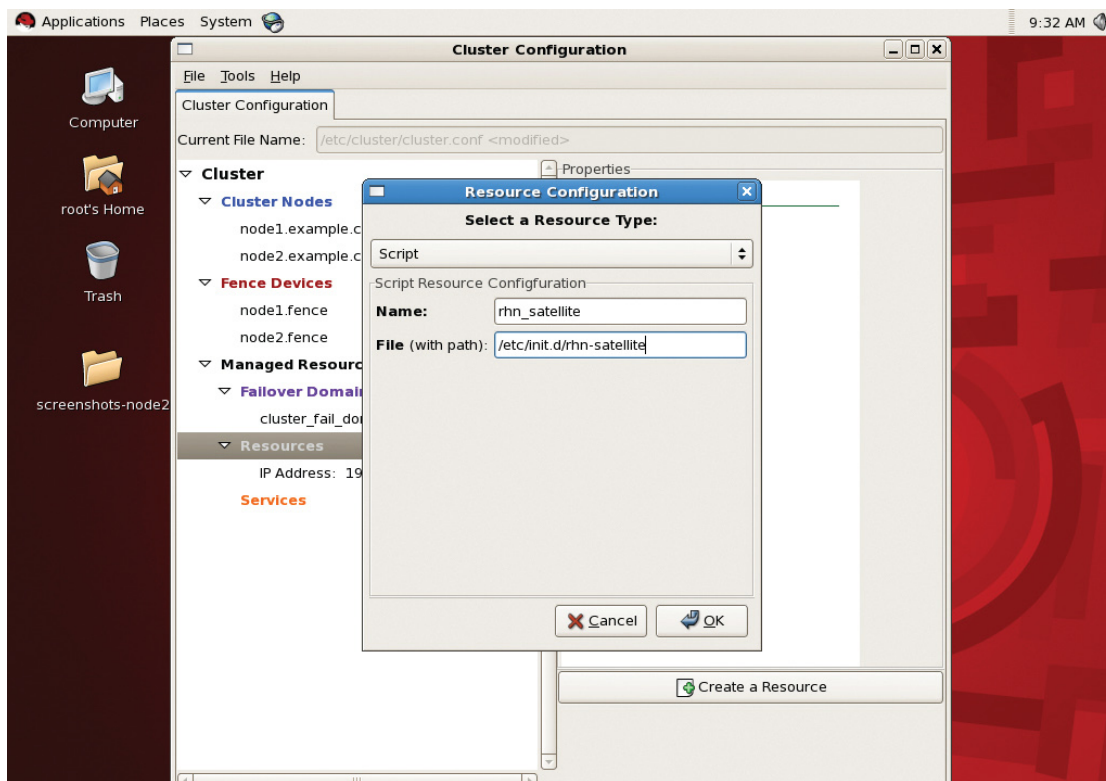
- In left pane, click "Failover Domains."
- In left pane, click "Create a Failover Domain."
- In new window, enter name of new failover domain.
- In new window, select nodes that are to be participants in failover domain from drop down menu.

Select node1 and node2

- If service is to be restricted to this group of machines, select "Restrict Failover To This Domain's Members" in right pane.
- If certain nodes in Failover domain are to have priority, assign priority by selecting "Prioritized List in right pane" and assign priorities by using "Adjust Priority" buttons in right pane.
- Click Close.
- Click File --> Save.
- Click OK in new window to save changes to cluster.conf file.



FIG 43: CREATE SCRIPT RESOURCE



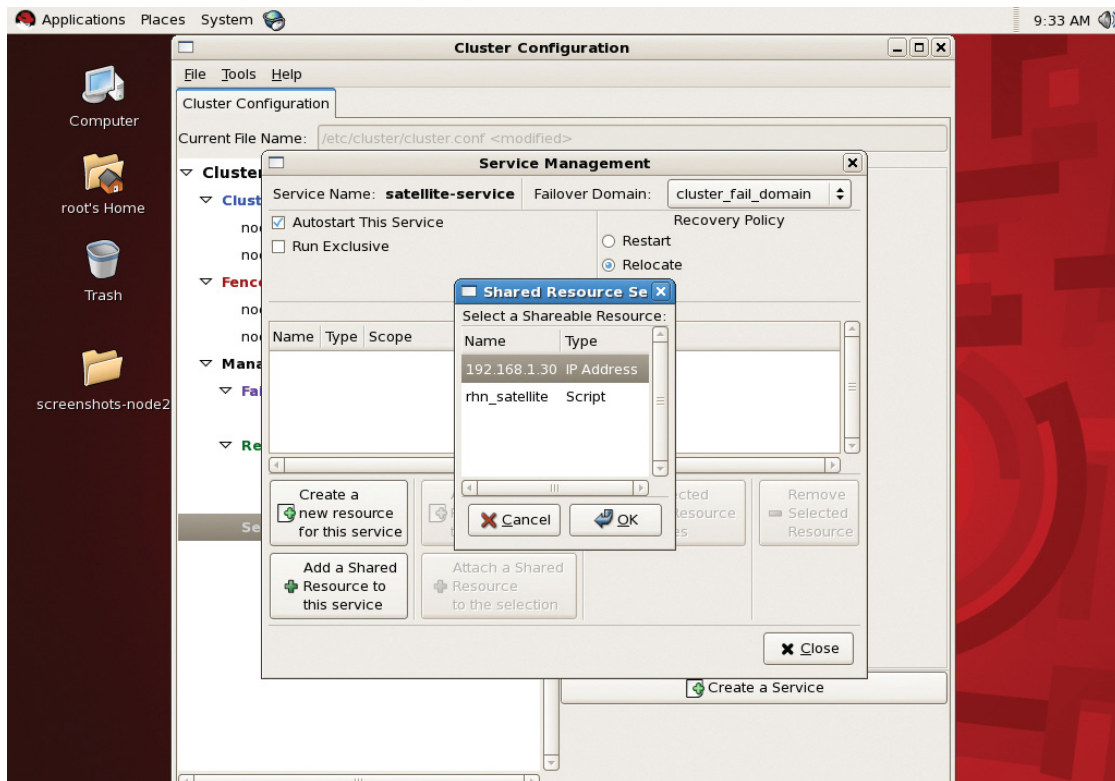
### Configure cluster resources

#### IP Address

- In left pane, click "Resources."
- In right pane, click "Create a resource."
- In new window, select "IP Address" from drop down menu.
- In new window, enter IP address.
- Click OK.



FIG 44A: CREATE SATELLITE SERVICE

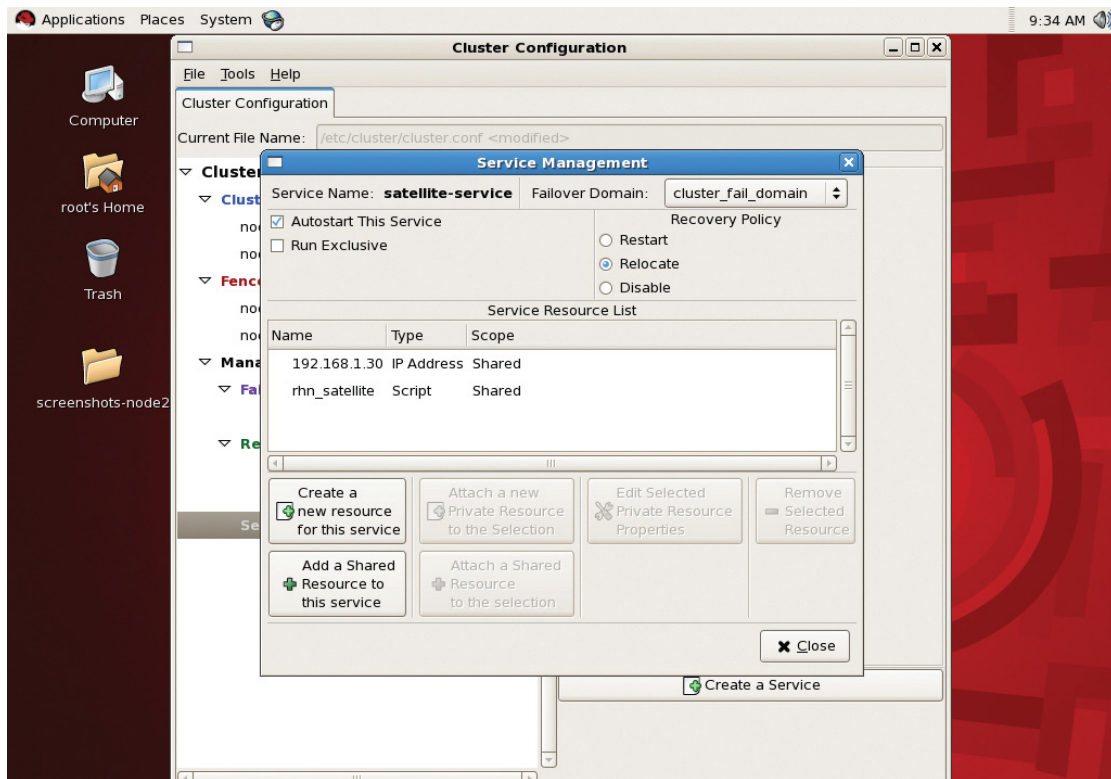


### Satellite script

- In left pane, click "Resources."
- In right pane, click "Create a resource."
- In new window, select "Script" from drop down menu.
- In new window, enter "rhnsatellite" for script name.
- Enter /etc/init.d/rhnsatellite in file section.
- Click Save.
- Click File --> Save.



FIG 45B: ASSIGN RESOURCES TO SERVICE

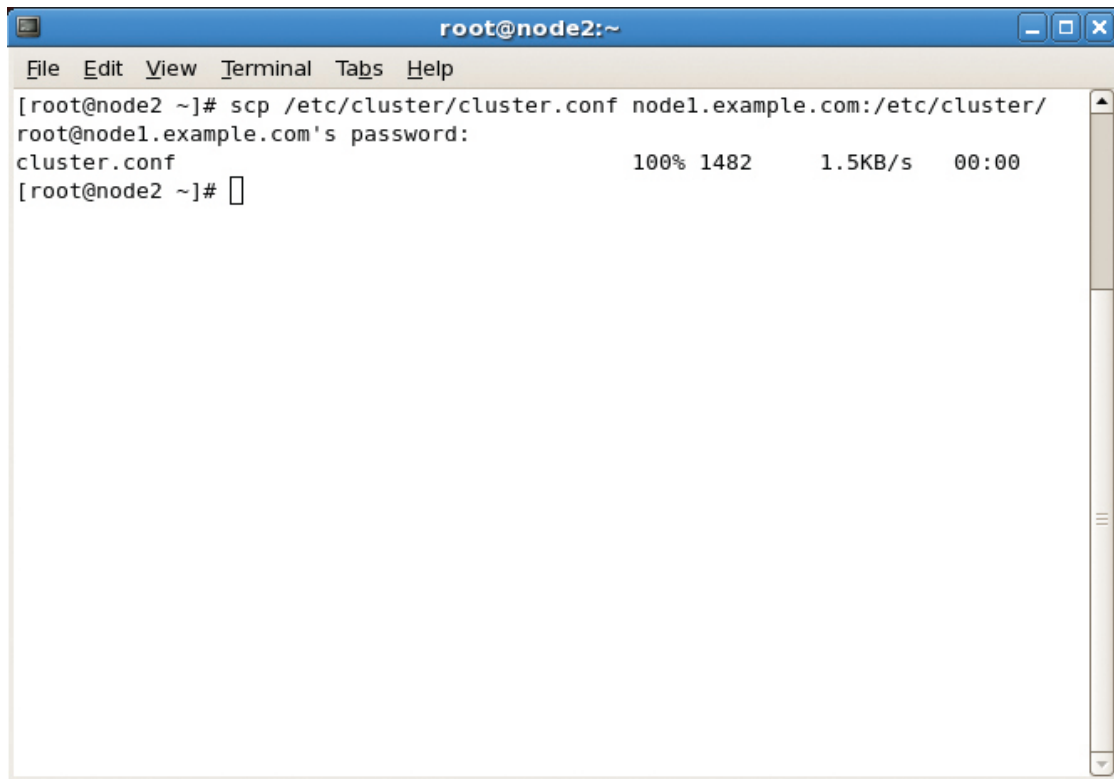


### Satellite service

- In left pane, click "Services."
- In right pane, click "Create a Service."
- In new window, enter name for new service.
- In top right of new window, select failover domain from drop down menu.
- In bottom left of window, click "Add a Shared Resource to this service."
- In new window, click "IP address shared resource," and click OK.
- In bottom left of window, click "Add a Shared Resource to this service."
- In new window, click "script name shared resource," and click OK.
- Click Close.
- Click File --> Save.



FIG 46: COPY CONFIGURATION FILE TO NODES



```
root@node2:~  
File Edit View Terminal Tabs Help  
[root@node2 ~]# scp /etc/cluster/cluster.conf node1.example.com:/etc/cluster/  
root@node1.example.com's password:  
cluster.conf 100% 1482 1.5KB/s 00:00  
[root@node2 ~]#
```

Close system-config-cluster application.

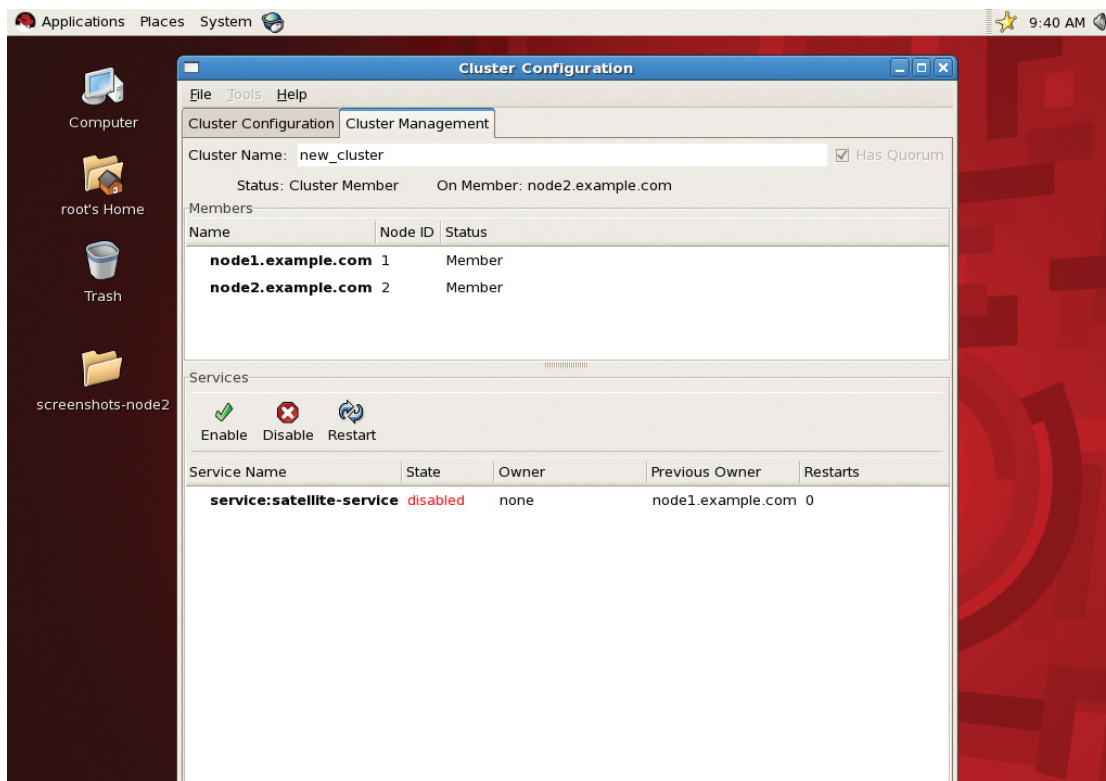
### ACTIVATE CLUSTER SUITE

To activate cluster suite, perform the following procedure:

```
# scp /etc/cluster/cluster.conf 192.168.1.2:/etc/cluster/
```



FIG 47: VERIFY CLUSTER OPERATION



From the command-line of both nodes:

```
# service cman start; chkconfig cman on
# service qdiskd start; chkconfig qdiskd on
# service rgmanager start; chkconfig rgmanager on
```


Start Cluster management tool and verify cluster management tab is available:

```
# system-config-cluster
```





FIG 48: START CLUSTER AND SERVICE



Start Red Hat Satellite on Cluster.

- Click management tab.
- In Service window, select "Red Hat Satellite service".
- Click enable button above.



**FIG 49: RELOCATE SATELLITE SERVICE**A terminal window titled 'root@node2:~' with a menu bar (File, Edit, View, Terminal, Tabs, Help). The terminal shows the command 'clusvcadm -r satellite-service -m node1.example.com' being executed. The output is: 'Trying to relocate service:satellite-service to node1.example.com...Success' followed by 'service:satellite-service is now running on node1.example.com'. The prompt returns to '[root@node2 ~]#'.

```
root@node2:~  
File Edit View Terminal Tabs Help  
[root@node2 ~]# clusvcadm -r satellite-service -m node1.example.com  
Trying to relocate service:satellite-service to node1.example.com...Success  
service:satellite-service is now running on node1.example.com  
[root@node2 ~]#
```

### TEST CLUSTER FAILOVER ABILITY

The following procedures were used to verify Satellite service performs failover properly.

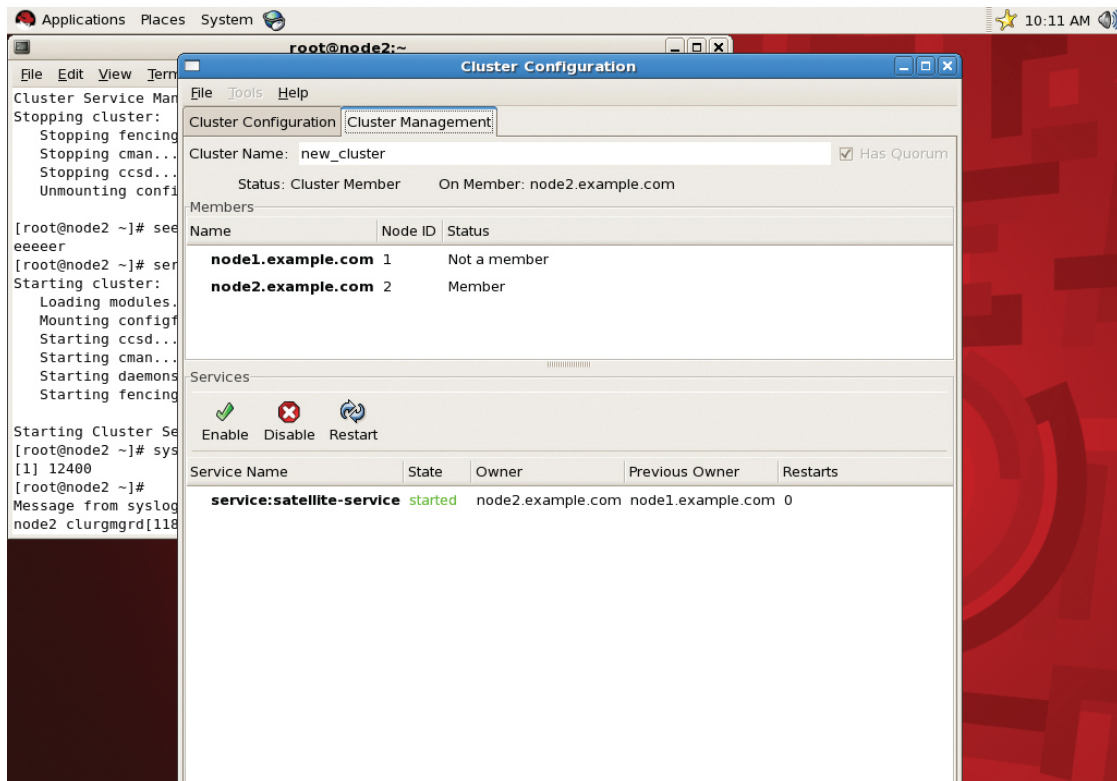
Manually relocate Satellite service:

1. On node currently running Satellite Service:

```
# clusvcadm -r rhn_satellite -m node2.example.com
```



FIG 50: VERIFY SATELLITE SERVICE RELOCATION



1. Observe cluster management window and verify relocation is successful.
2. While Satellite service is running, power off cluster node.

```
# poweroff
```

2. In Cluster management window, verify cluster service relocates successfully.

Disconnect ethernet device:

1. While satellite service is active, manually disconnect ethernet devices on cluster node currently running satellite service.
2. In Cluster management window, verify service is successfully relocated.



Block ping (this will test heuristics).

1. If node 1 is running the service, performing the following on node 1:

```
# tail -f /var/log/messages
```

2. On node 2 perform the following:

```
# iptables -A OUTPUT -d 192.168.1.30 -j REJECT
```

This will cause node 2 to fail the heuristics test and be declared dead to the cluster manager, which will cause it to be fenced.

## RECOMMENDATIONS

---

Configure cron job backing up Oracle embedded database. It is recommended that database backups be performed at a regular interval to maintain integrity. It is advisable to schedule a cron job that performs a database backup on a nightly basis for immediate recovery.

File system check of shared resources. Due to the shared resource mount points being employed, it is often recommended that file system checks are performed on LVM partitions.

## CONCLUSION

---

Combining the high availability features of Red Hat Cluster Suite with Red Hat Network Satellite provides environments with a complete systems management tool with maximum up-time. The above procedure demonstrates how Red Hat technologies can be combined and implemented in enterprise environment to meet specific needs.

### RED HAT ENTERPRISE LINUX

Red Hat Enterprise Linux 5 Release Notes

→<http://www.redhat.com/docs/manuals/enterprise/#RHEL5>

Red Hat Enterprise Linux 5 Installation Guide

→[http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5/html/Installation\\_Guide/index.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/index.html)

Red Hat Enterprise Linux 5 Deployment Guide

→[http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5/html/Deployment\\_Guide/index.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Deployment_Guide/index.html)

### RED HAT NETWORK SATELLITE

Red Hat Network Satellite Release Notes

→[http://www.redhat.com/docs/manuals/satellite/Red\\_Hat\\_Network\\_Satellite-5.2.0/html/Release\\_notes/index.html](http://www.redhat.com/docs/manuals/satellite/Red_Hat_Network_Satellite-5.2.0/html/Release_notes/index.html)



## Red Hat Network Satellite Installation Guide

→[http://www.redhat.com/docs/manuals/satellite/Red\\_Hat\\_Network\\_Satellite-5.2.0/html/Installation\\_Guide/index.html](http://www.redhat.com/docs/manuals/satellite/Red_Hat_Network_Satellite-5.2.0/html/Installation_Guide/index.html)

## Red Hat Network Satellite Client Configuration Guide

→[http://www.redhat.com/docs/manuals/satellite/Red\\_Hat\\_Network\\_Satellite-5.2.0/html/Client\\_Configuration\\_Guide/index.html](http://www.redhat.com/docs/manuals/satellite/Red_Hat_Network_Satellite-5.2.0/html/Client_Configuration_Guide/index.html)

## RED HAT CLUSTER SUITE

### Red Hat Cluster Suite Overview for Red Hat Enterprise Linux 5.2

→[http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5.2/html/Cluster\\_Suite\\_Overview/index.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Cluster_Suite_Overview/index.html)

### Red Hat Cluster Suite Configuration and Management

→[http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5.2/html/Cluster\\_Administration/index.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Cluster_Administration/index.html)

## APPENDIX

Cluster.conf

```
<?xml version="1.0" ?>
<cluster config_version="6" name="new_cluster">
  <quorumd device="/dev/sdd1" interval="2" label="quorum" min_score="1" tko="10"
  votes="1">
    <heuristic interval="2" program="ping -c1 -t1 192.168.1.1" score="1"/>
  </quorumd>
  <fence_daemon post_fail_delay="0" post_join_delay="3"/>
  <clusternodes>
    <clusternode name="node1.example.com" nodeid="1" votes="1">
      <fence>
        <method name="1">
          <device name="node1.fence"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="node2.example.com" nodeid="2" votes="1">
      <fence/>
    </clusternode>
  </clusternodes>
  <cman expected_votes="3"/>
</cluster>
```



```
<fencedevices>
  <fencedevice agent="fence_rsa" ipaddr="192.168.1.11" login="user" name="node1.
fence" passwd="password"/>
  <fencedevice agent="fence_rsa" ipaddr="192.168.1.21" login="user" name="node2.
fence" passwd="password"/>
</fencedevices>
<rm>
  <failoverdomains>
    <failoverdomain name="cluster_fail_domain" ordered="1" restricted="1">
      <failoverdomainnode name="node1.example.com" priority="1"/>
      <failoverdomainnode name="node2.example.com" priority="2"/>
    </failoverdomain>
  </failoverdomains>
  <resources>
    <ip address="192.168.1.30" monitor_link="1"/>
    <script file="/etc/init.d/rhn-satellite" name="rhn_satellite"/>
  </resources>
  <service autostart="1" domain="cluster_fail_domain" name="satellite-service"
recovery="relocate">
    <ip ref="192.168.1.30"/>
    <script ref="rhn_satellite"/>
  </service>
</rm>
</cluster>
```



## LEARN MORE ABOUT RED HAT CONSULTING

---

[www.redhat.com/consulting](http://www.redhat.com/consulting)

### EUROPE, MIDDLE EAST AND AFRICA

00800 7334 2835  
[www.europe.redhat.com](http://www.europe.redhat.com)  
[europe@redhat.com](mailto:europe@redhat.com)

**Turkey:** 00800 448 820 640

**Israel:** 1809 449 548

**UAE:** 80004449549

### ASIA PACIFIC

+65 6490 4200  
[www.apac.redhat.com](http://www.apac.redhat.com)  
[apac@redhat.com](mailto:apac@redhat.com)

**ASEAN:** 800 448 1430

**Australia and New Zealand:**

1800 733 428

**Greater China:** 800 810 2100

**India:** +91 22 3987 8888

**Japan:** 0120 266 086

**Korea:** 080 708 0880

### NORTH AMERICA

1-888-REDHAT1  
[www.redhat.com](http://www.redhat.com)

### LATIN AMERICA

+54 11 4341 6200  
[www.latam.redhat.com](http://www.latam.redhat.com)  
[info-latam@redhat.com](mailto:info-latam@redhat.com)