



Setting Up Prime Central High Availability in a Bare-Metal Environment

Revised: April 4, 2014, OL-28573-02

Installation Overview

The Prime Central RHCS local redundancy HA configuration in a Bare-metal environment has the following characteristics and requirements:

- Prime Central and the embedded database are installed in a dual-node cluster. [Figure 1-1](#) shows a basic dual-node, local redundancy cluster.
- RHCS must be installed on both cluster nodes. Each node has the platform to run both Prime Central and the database services. RHCS manages the local HA setup.
- RHCS requires a disk resource that is mountable from both nodes.
- The Prime Central installer places Prime Central and the database on the node where you ran the installation. After installation, you can relocate them as needed.
- Prime Central, the Integration Layer, and the database services must be placed on the same server.
- Prime Central and the embedded database are always mounted with external shared storage.
- Prime Central does not recognize RHCS. RHCS continuously obtains the cluster status by running a set of scripts. If a problem occurs, RHCS unmounts, then remounts the appropriate server and database. Therefore, every node in the HA setup must be able to mount the storage.


Note the following:

- The Prime Central and oracle home directories must be created manually under the mounted storage. This ensures that the operating system (OS) user created on both servers has a home directory available, even if the storage is moved to another node. These directories must have relevant permissions for the network user and oracle user.
- Each service has its own virtual IP address (virtual IP). This means Prime Central clients treat a failover or switchover like a local service restart.
- Only one instance of the Prime Central files exists, and it is located on the shared storage. Duplicate user and home directories are created on each node as placeholders. If a switchover occurs, the storage unmounts from one node and mounts on the other.

- A local redundancy configuration requires a *fencing* hardware unit for cutting off a node from the shared storage. Fencing ensures data integrity and prevents a “split-brain scenario” where servers are disconnected from each other, and each assumes the other server failed. If a failure occurs, the cutoff can be performed by:
 - Powering off the node with a remote power switch
 - Disabling a switch port fiber channel
 - Revoking a host’s SCSI 3 reservations

If a problem with the cluster node occurs, RHCS invokes the fencing device with the peer and waits for the success signal.

Figure 1-1 Prime Central Dual-Node, Local Redundancy Cluster



1	External storage	3	Local cluster node 2
2	Local cluster node 1		

RHCS Components and Functionality

RHCS is included with the Red Hat Enterprise Linux 5.8/6.5 (RHEL 5.8/6.5) Advanced Program and has the following components:

- Cluster infrastructure—Basic functions that enable a group of computers (nodes) to work together as a cluster. The RHCS cluster infrastructure performs cluster management, lock management, fencing, and cluster configuration management.
- High availability service management—Failover of services from one cluster node to another when a node becomes inoperative.
- Cluster administration tools—Configuration and management tools for setting up, configuring, and managing a Red Hat cluster, including cluster infrastructure components, high availability service management components, and storage.

Supported RHCS Fencing Options

A fencing device cuts off a node from shared storage to ensure data integrity. The supported fencing options are:

- *fence_ilo*—Cisco validated Prime Central HA with the HP Integrated Lights-Out (HP iLO) *fence_ilo* automatic fencing method; other automatic fencing methods can be used but have not been validated.
- *fence_manual*—Assigns the manual fencing agent. This fencing agent is temporary; you should not use it in production because it does not perform automated actions. If a cluster problem occurs, you must manually disconnect the node and storage, or use another fencing agent to disconnect them. If you choose this option during the installation because you want to add a different Red Hat-supported fencing device, use the RHCS GUI to provision the device after installation. Be sure to add it as the main fencing method, and move the manual fencing agent to the backup method, as shown in [Figure 1-2](#) and [Figure 1-3](#).


Note the following:

- In [Figure 1-2](#), the HP iLO host 1 is 192.168.1.10.
- In [Figure 1-3](#), the HP iLO host 2 is 192.168.1.11.
- To prevent fencing loops, the cluster interconnect and power fencing (for example, HP iLO) should use the same network, such as bond0.
- If the main fencing device is a remote power switch, define all ports that simultaneously supply power to the node.

**Note**

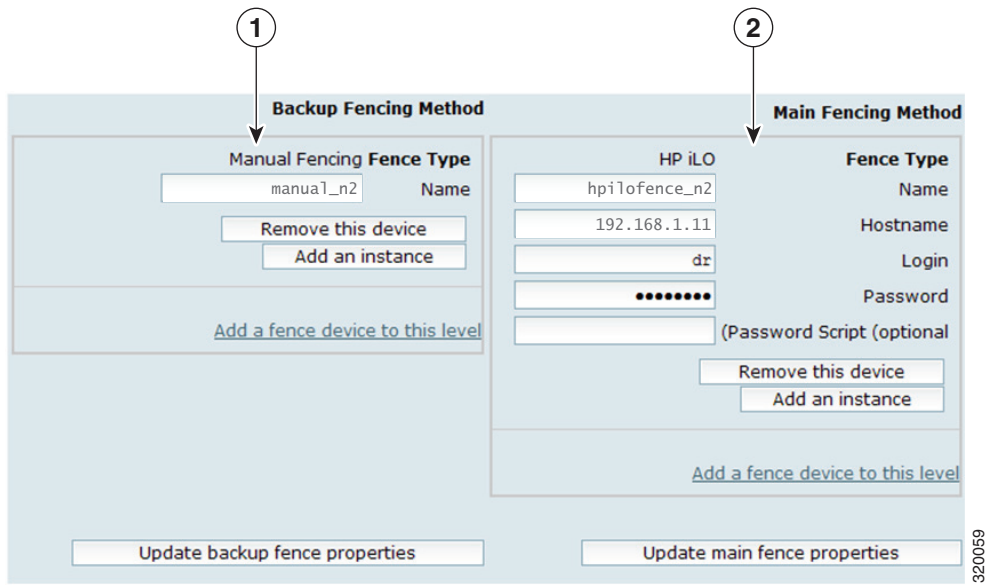
For general information about the RHCS web GUI, see [Verifying the Prime Central RHCS HA Installation, page 1-9](#). For complete information about using the RHCS web GUI, see the *Red Hat Conga User Guide*. For information about provisioning and managing dual-node cluster fencing devices, see the RHCS user documentation.

Figure 1-2 RHCS GUI Fencing Method Window—Node 1



1	Backup fencing method	2	Main fencing method
----------	-----------------------	----------	---------------------

Figure 1-3 RHCS GUI Fencing Method Window—Node 2



1	Backup fencing method	2	Main fencing method
----------	-----------------------	----------	---------------------

The fencing methods shown in [Figure 1-2](#) and [Figure 1-3](#) change the following sections in the `cluster.conf` file:

```
<cluster>
<clusternodes>
  <clusternode name="prime-central-linux1.cisco.com" nodeid="1" votes="1">
    <multicast addr="224.0.0.251" interface="eth0"/>
    <fence>
      <method name="method_1">
        <device name="hпилofence1"/>
      </method>
    </fence>
  </clusternode>
  <clusternode name="prime-central-linux2.cisco.com" nodeid="2" votes="1">
    <multicast addr="224.0.0.251" interface="eth0"/>
    <fence>
      <method name="method_2">
        <device name="hпилofence2"/>
      </method>
    </fence>
  </clusternode>
  ...
<fencedevices>
  <fencedevice agent="fence_ilo" hostname="value-of-HP-iLO-host-1" login="prime"
  name="hпилofence1" passwd="password"/>
  <fencedevice agent="fence_ilo" hostname="value-of-HP-iLO-host-2" login="prime"
  name="hпилofence2" passwd="password"/>
</fencedevices>
  ...
</cluster>
```

RHCS HA Local Redundancy Requirements


Table 1-1 Prime Central RHCS HA Local Redundancy Requirements

Area	Requirement
OS	Red Hat Enterprise Linux 5.8/6.5 with the Red Hat Clustering Suite. Note We recommend that customers check periodically for RHEL 5.8/6.5 patches and install any available updates.
Oracle	11gR2 Enterprise Edition. Note Oracle 11gR2 EE is included in the Prime Central embedded database installation.
Hardware	RHEL 5.8/6.5-certified platform with fencing capabilities.
Network	<ul style="list-style-type: none"> Cluster nodes must be able to communicate with each other using multicast. Each network switch and the associated networking equipment in a Red Hat cluster must be configured to enable multicast addresses and support IGMP. Without multicast and IGMP, not all nodes can participate in a cluster, causing the cluster to fail. Network timing must be configured. Note During the RHCS HA installation, you will be asked to confirm that NTP timing is configured.

Table 1-1 Prime Central RHCS HA Local Redundancy Requirements (continued)

Area	Requirement
Storage and file system	RHCS requires shared Fibre Channel, Fibre Channel over Ethernet (FCoE), or Internet Small Computer System Interface (iSCSI) storage that presents an ext3 file system and is accessible from all cluster nodes.
Disk space	5 GB under /tmp.
Miscellaneous	<ul style="list-style-type: none"> The same subnet must be assigned to all nodes, including all services (virtual IP addresses in the same subnet). All nodes must have RHEL 5.8/6.5 installed. All systems must be homogeneous with the same configuration versions and packages. Shared storage must not be auto-mounted, because RHCS performs the mounting. Use one partition for each cluster service. For a single shared disk, use a single partition for each service on the same disk. In other words, the shared storage must not appear in /etc/fstab. All shared storage units must be configured with a label, which RHCS uses to mount and unmount storage. Virtual IP addresses must be assigned for each service. IP addresses assigned to services should not be attached to any server. RHCS will manage them; that is, it will add and remove them from the server that is running the service. Fencing devices must be deployed. Multicast communication must not be blocked by a firewall.

We strongly recommend that you use a hardware installation designed to avoid a single point of failure. See [Figure 1-4](#).

Figure 1-4 Local Redundancy Hardware Installation Designed to Avoid a Single Point of Failure


246599

1	Dual NICs	4	NIC bonding (active/backup mode)
2	Disk mirroring	5	Each NIC connects to a separate switch
3	Redundant RAID controllers		

Configuring Hardware for Prime Central RHCS Local Redundancy HA

Figure 1-5 shows the recommended hardware configuration for Prime Central RHCS local redundancy HA.

Figure 1-5 Prime Central Dual-Node Cluster for Local Redundancy HA



1	Prime Central external storage: one Prime Central volume	4	Oracle external storage: One to three data volumes One archive volume Zero to three redo log volumes
2	Prime Central server: Two internal disks One OS One mirror	5	Oracle database server: Two internal disks One OS One mirror
3	Dual Gigabit Ethernet crossover connections	6	Dual connections from each server to each external disk storage unit

Configure the external storage so all disks and logical unit numbers (LUNs) are accessible from both servers in the cluster. The disk and LUN configuration depends on the storage type:

- If you are using JBOD disks, provide enough physical disks to create the volumes listed in [Table 1-2](#) to satisfy the Oracle performance requirements.
- If you are using storage that supports hardware RAID, divide the physical disks into LUNs so that the volumes listed in [Table 1-2](#) can be created and configured to satisfy the Oracle performance requirements and can be protected with RAID5, RAID1, or RAID10. The Oracle volumes can be created on a single LUN.

Table 1-2 Volume Sizes

Volume	Minimum Size	Comments
Prime Central	50 GB	—
Oracle application + data files	10 GB	—
Oracle redo logs	12.8 GB	—
Oracle archives	20 GB	Contact your Cisco account representative for information.
Oracle additional data files (if used)	—	Based on Prime Central alarm history.
Oracle backup	50 GB	—

Configuring the OS for Prime Central Local Redundancy HA Managed by RHCS

1. Install the OS and all recommended patches on both servers in the cluster. Installations on both servers must be identical.
2. Verify that access is available to all external disks.
3. Create the internal disk partitions listed in [Table 1-3](#). Placing the individual directories in separate partitions is recommended, but not required.
4. Complete the internal disk partitions for both servers.
5. Keep the nodes synchronized:

```
# echo server tick.redhat.com$\n'restrict tick.redhat.com mask 255.255.255.255
nomodify notrap noquery >> /etc/ntp.conf
# chkconfig ntpd on
# service ntpd start
```

Table 1-3 Disk Partitions

Partition	Space Required (MB)
swap	Standard amount of space, as per the system configuration
/tmp	Standard amount of space + 5120
/	Standard amount of space + 6144
/var	Standard amount of space + 1024 for HA utilities
/usr/local/bin	Standard amount of space + 200 for cluster utilities
/etc	Standard amount of space + 200 for the cluster configuration

Installing the Red Hat Cluster Service

Using the procedures in the Red Hat user documentation, install RHEL 5.8/6.5 with the RHCS. When you set up the RHCS disk groups and volumes, keep the following in mind:

- All of the shared storage should have an ext3 file system installed and a label set.
- Shared storage must be accessible from all cluster nodes.

Verifying the Prime Central RHCS HA Installation

Table 1-4 Local Redundancy Verification Tests

Description	Procedure	Expected Results
Local Cluster Hardware Failure		
<p>Name: Cluster node hardware failure.</p> <p>Purpose: Test the local site failover (including fence test) due to node failure.</p>	<ol style="list-style-type: none"> 1. Aggressively power off the active node that runs both services (Prime Central and DB). 2. Verify that both services are relocated to the redundant node. 	<p>Within several minutes, the redundant cluster node identifies that the active node is not available and fences it, evicting it from the cluster and relocating all the services to the only remaining node.</p>
Manual Cluster Administration		
<p>Name: Manual service stop.</p> <p>Purpose: Verify that the service can be stopped manually.</p>	<ol style="list-style-type: none"> 1. Enter clusvcadm -d <i>service-name</i>. 2. Verify that the service is not running and no errors appear in the cluster log (/var/log/messages for both cluster nodes). 	<p>The stopped service is no longer running.</p>
<p>Name: Manual service start.</p> <p>Purpose: Verify that the service can be started manually.</p>	<ol style="list-style-type: none"> 1. Enter clusvcadm -e <i>service-name</i>. 2. Verify that the service is running and no errors exist in the cluster log (/var/log/messages on both cluster nodes). 	<p>The service is running.</p>
<p>Name: Manual service restart.</p> <p>Purpose: Verify that the service can be restarted manually.</p>	<ol style="list-style-type: none"> 1. Enter clusvcadm -R <i>service-name</i>. 2. Verify that the service is running and no errors exist in the cluster log (/var/log/messages on both cluster nodes). 	<p>The service is running.</p>
<p>Name: Manual service relocation.</p> <p>Purpose: Verify that the service can be relocated manually.</p>	<ol style="list-style-type: none"> 1. Enter clusvcadm -r <i>service-name</i>. 2. Verify that the service is not running on the current node and is running on the standby node. 3. Verify that no errors appear in the cluster log (/var/log/messages on both cluster nodes). The service is stopped on the active node and then started on the redundant node. 4. Test both the Prime Central and Oracle services. 	<p>The service is stopped on the active node and started on the redundant node.</p>

Table 1-4 Local Redundancy Verification Tests (continued)

Description	Procedure	Expected Results
Ordered Cluster Node Startups		
<p>Name: Node startup in existing cluster.</p> <p>Purpose: Verify that a cluster node starts up and rejoins a cluster after it is restarted.</p>	<ol style="list-style-type: none"> Restart one of the cluster nodes. Verify that the node joins the cluster after the reboot. Relocate one of the services to the rebooted node and verify that it is running. Check the log for errors. 	The rebooted node joins the cluster and runs the services.
<p>Name: Simultaneous node startup.</p> <p>Purpose: Verify that the cluster is set up correctly when both nodes start simultaneously.</p>	<ol style="list-style-type: none"> Start both nodes from the power off state. Verify that both nodes appear in the cluster after they are up, with both services running on the cluster. Check the log for errors. 	Both cluster nodes join the cluster; both services are running.
<p>Name: Single-node startup.</p> <p>Purpose: Test the cluster functionality when only one is node running.</p>	<ol style="list-style-type: none"> Power down both nodes, then start one of them. The running node fences the other node and runs the services. The fenced node joins the cluster to create the dual-node cluster. Check the log for errors. 	Both cluster nodes join the cluster; both services are running.
Local Cluster Service Failure		
<p>Name: Service failure.</p> <p>Purpose: Test the service startup after a failure occurs.</p>	<ol style="list-style-type: none"> Simulate a service failure by stopping its processes or shutting down the Oracle listener. Verify that the service restarts on the same node where it was running. Check the log for errors. Test both the Prime Central and Oracle services. 	The service is restarted on the same node.

Table 1-4 Local Redundancy Verification Tests (continued)

Description	Procedure	Expected Results
Local Cluster Hardware Failure		
Name: Stop node with fencing off. Purpose: Verify that the node requires manual fencing after the other node, including its fencing agent, is removed.	<ol style="list-style-type: none"> 1. Disconnect the fencing agent from one of the nodes, then power it off. 2. Observe the other node behavior. 3. Check the log for errors and the request for manual fencing. 	The fence_ack_manual required notification appears in the log files. The cluster is running with one node and all services running on it.
Name: Single-node cluster. Purpose: Verify that the cluster can function when the other node does not exist or has no power.	<ol style="list-style-type: none"> 1. Power down both nodes. 2. Disconnect the fencing agent from one of the nodes. 3. Start the other node. It attempts to fence the other node, but fails with the regular fencing agent. Manual fencing is required. 4. Acknowledge the manual fencing. 	The cluster does not start the services (and does not show in the clustat command) before acknowledging that manual fencing is performed.

Installing Prime Central and an Embedded Database in a Local Redundancy HA Configuration

Installing Prime Central in an RHCS local HA configuration is a three-part process:

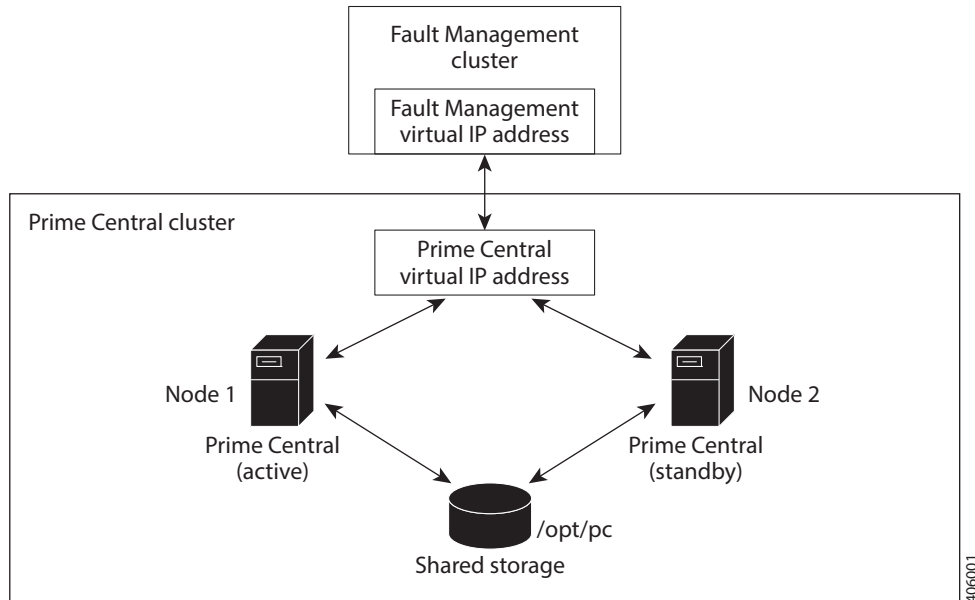
1. Install RHEL 5.8/6.5 on both nodes.
2. Use multipath shared storage and install Prime Central on node 1.
3. Configure and enable clustering so that Prime Central can relocate between nodes.

The examples provided use the following hostnames and IP addresses; yours will be different:

- Node 1—prime-ha-node1.cisco.com (192.168.1.110)
- Node 2—prime-ha-node2.cisco.com (192.168.1.120)
- Virtual IP address—prime-service.cisco.com (192.168.1.130)
- Gateway—192.168.1.1
- Domain Name System (DNS)—192.168.1.2

Figure 1-6 shows an example of a Prime Central cluster in an HA configuration.

Figure 1-6 Prime Central Cluster in an HA Configuration



Before You Begin

- Verify that your system meets all the hardware and software requirements in “Installation Requirements” in the [Cisco Prime Central 1.4.1 Quick Start Guide](#).
- Set up two nodes that have:
 - Static IP addresses and hostnames that are registered correctly in the DNS.
 - The same root password, which cannot contain a percent sign (%), ^, \$, *).
- Set up one virtual IP address and hostname that are registered correctly in the DNS. In this section, the virtual IP address is 192.168.1.130.
- Set up shared storage that is compatible with RHEL device-mapper (DM) multipath and cluster fencing.
- Install RHEL 5.8/6.5 on both nodes.
- If the default folder location and names are changed, look for the section “Require Manual Definition” to make corresponding changes in the below mentioned files:
 - /root/ha-stuff/pc/PrimeCentral.sh
 - /root/ha-stuff/pc/UninstallPrimeCentral.sh
 - /usr/local/bin/pc.sh
 - /usr/local/bin/createUserGroup.sh

Adding Clustering to the Installed Red Hat Server

To add clustering to the newly installed Red Hat server, complete the following steps in parallel on both nodes, except where noted:

Step 1 Create local directories named `/rhel` and `/cdrom`.

Step 2 Copy the `.iso` file that was used for the virtual machine (VM) RHCS installation to the `/rhel` directory.

Step 3 Mount the `/rhel .iso` file to `/cdrom`:

```
# cd /rhel
# mount -t iso9660 -o loop /rhel/rhel-server-5.8-x86_64-dvd.iso /cdrom
or
# mount -t iso9660 -o loop /rhel/rhel-server-6.5-x86_64-dvd.iso /cdrom
```



Note To permanently mount the drive, update the `/etc/fstab` file. See http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/4/html/Introduction_To_System_Administration/s2-storage-mount-fstab.html.

Step 4 Create a file named `/etc/yum.repos.d/local.repo`. Use UNIX format and be sure there are no spaces before lines.

Step 5 Save the newly created file in `local.repo`, as follows:

```
[local]
name=Red Hat Enterprise Linux $releasever - $basearch - Local
baseurl=file:///cdrom/Server
enabled=1
gpgcheck=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
[cluster]
name=Red Hat Enterprise Linux $releasever - $basearch - Cluster
baseurl=file:///cdrom/Cluster
enabled=1
gpgcheck=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
```

Step 6 Install the clustering package:

```
# yum groupinstall Clustering
```

Step 7 Add the information for both nodes to the `/etc/hosts` file; for example:

```
192.168.1.110    prime-ha-node1.cisco.com    prime-ha-node1
192.168.1.120    prime-ha-node2.cisco.com    prime-ha-node2
```

Step 8 Generate a Secure Shell (SSH) key for the root user:

```
# chmod 755 ~
# ssh-keygen -t rsa -N "" -b 2047 -f ~/.ssh/id_rsa
# chmod 600 ~/.ssh/id_rsa
```

Step 9 (On the first node only) Share the node's public key with the other node so that dynamically creating a secure shell between the nodes does not prompt for a password:

```
# rsync -av ~/.ssh/id_rsa.pub ~/.ssh/authorized_keys
# ssh root@node2 "cat ~/.ssh/id_rsa.pub" >> ~/.ssh/authorized_keys
# rsync -av ~/.ssh/authorized_keys root@prime-ha-node2.cisco.com:/root/.ssh/
```

Step 10 Verify that the `.ssh` directory has 700 permission and the `.ssh/id_rsa` file has 600 permission:

```
# chmod 700 ~/.ssh
# chmod 600 ~/.ssh/id_rsa
```

Step 11 Verify that your SSH is working without an authentication or password prompt:



Caution

The Prime Central service will not start if SSH prompts for authentication or a password. Be sure to complete all of the following substeps.

a. On node `prime-ha-node1.cisco.com`, enter:

```
# ssh root@prime-ha-node2.cisco.com
# exit
# ssh root@prime-ha-node2
# exit
# ssh root@192.168.1.120
# exit
```

b. On node `prime-ha-node2.cisco.com`, enter:

```
# ssh root@prime-ha-node1.cisco.com
# exit
# ssh root@prime-ha-node1
# exit
# ssh root@192.168.1.110
# exit
```

c. If you are prompted for a password, check the permissions of all folders and files that you modified in the preceding steps.

d. If you are prompted to continue connecting, enter **yes**. (The prompt should appear only the first time you use SSH to connect to the node.)

Step 12 Verify that the virtual IP address is accessible from outside the cluster's subnet:

```
# ip addr add 192.168.1.130 dev eth0
```

Step 13 On a computer outside the cluster's subnet, ping the virtual IP address:

```
# ping 192.168.1.130
# ip addr del 192.168.1.130 dev eth0
```

If you do not get a valid response, determine which part of the OS or network setup is blocking.

Adding Shared Partitions

To add shared partitions, complete the following steps in parallel on both nodes, except where noted:


Note

The examples provided use device mapping names such as `mpath2` and `mpath2p1`; yours will be different.

Step 1 Set up multipath:

```
# vi /etc/multipath.conf
-- Comment out 'blacklist' section
-- For example:
-- #blacklist {
--     # devnode "*"
-- }
# modprobe dm-multipath
# service multipathd start
# chkconfig multipathd on
```

Step 2 Check for available disks:

```
cd /dev/mapper/
ls -la
total 0
drwxr-xr-x  2 root root    120 May  4 18:42 .
drwxr-xr-x 13 root root   3940 May  4 18:42 ..
crw-----  1 root root    10, 63 May  4 18:42 control
brw-rw----  1 root disk 253,  2 May  4 18:42 mpath2
brw-rw----  1 root disk 253,  0 May  4 18:42 VolGroup00-LogVol00
brw-rw----  1 root disk 253,  1 May  4 18:42 VolGroup00-LogVol01
```

In the output, note `mpath2`, which is the multipath virtual device or disk that you will use later as shared storage.


Note

If you previously set up a partition on the disk, you might see output such as `mpath2p`. You must delete that partition before proceeding to the next step.

Step 3 (On the first node only) Create a 100-GB, shared partition:

```
# fdisk mpath2
Command (m for help): p
Command (m for help): n
Command action
  e   extended
  p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-19581, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-19581, default 19581): +100GB
Command (m for help): w
```

Step 4 Reboot both nodes.

Step 5 Check for new partitions:

```
# cd /dev/mapper/
# ls -la
total 0
drwxr-xr-x  2 root root    180 May  4 19:49 .
drwxr-xr-x 13 root root   4120 May  4 19:49 ..
crw-----  1 root root    10, 63 May  4 19:49 control
brw-rw----  1 root disk 253,  2 May  4 19:49 mpath2
brw-rw----  1 root disk 253,  3 May  4 19:49 mpath2p1
brw-rw----  1 root disk 253,  0 May  4 19:49 VolGroup00-LogVol100
brw-rw----  1 root disk 253,  1 May  4 19:49 VolGroup00-LogVol101
```

Step 6 (On the first node only) Format the new shared partition:

```
# mkfs.ext3 /dev/mapper/mpath2p1
```

Step 7 Create target locations:

```
# mkdir /opt/pc
```

Step 8 Verify that both nodes can mount and unmount the shared storage:

- a. On the first node, mount the shared storage and save a file that contains only the value *1* to the shared storage. The test.txt file should exist in the list of contents of /opt/pc:

```
# mount /dev/mapper/mpath2p1 /opt/pc
# vi /opt/pc/test.txt
1
:wq
# ls -la /opt/pc
# umount /opt/pc
```

- b. On the second node, mount the shared storage and verify that the test.txt file exists and contains the value *1*:

```
# mount /dev/mapper/mpath2p1 /opt/pc
# vi /opt/pc/test.txt
:q
# umount /opt/pc
```

If you cannot mount or unmount the shared storage, or if the test.txt file does not exist when you mount it to the second node, your multipath is not set up correctly.

Accessing and Distributing the Prime Central .tar File

Step 1 Insert the Cisco Prime Central 1.4.1 USB, navigate to the High Availability/RHCS Bare Metal Local HA/Prime Central folder, and locate the primecentral_v1.4.1_ha_vm.tar.gz file.

Step 2 Use SSH to connect to the first node.

Step 3 Copy the primecentral_v1.4.1_ha_vm.tar.gz file to the first node.

Step 4 Back up the following directories on both nodes:

- /root/ha-stuff/pc
- /usr/local/bin

Step 5 Distribute the file:

```
# tar -zxf primecentral_v1.4.1_ha_vm.tar.gz -C / --owner root --no-same-owner
```

Step 6 Navigate to the Base Application folder and copy primecentral_v1.4.1.bin and all available zip files into the /root/ha-stuff/pc directory:

```
# chmod 755 /usr/local/bin/*  
# chmod 755 /root/ha-stuff/pc/*
```

Installing Prime Central on the First Node in an HA Setup

To install Prime Central on the first node only:

Step 1 Mount the shared partitions:

```
# mount /dev/mapper/mpath2p1 /opt/pc
```

Step 2 Add a virtual IP cluster service address for the Prime Central service:

```
# ip addr add 192.168.1.130 dev eth0
```

Step 3 Update the install.properties file and verify that all required properties have values. Review the comments at the top of the install.properties file for details.



Note To install Prime Central silently, you must edit the /root/ha-stuff/pc/install.properties file. See “Sample install.properties Files” in the [Cisco Prime Central 1.4.1 Quick Start Guide](#).

Step 4 Install Prime Central:

```
# cd /root/ha-stuff/pc  
# ./PrimeCentral.sh 192.168.1.130 node's-root-password second-node-IP-address
```



Note You run the PrimeCentral.sh script by adding the preceding command-line parameters. If you do not add the command-line parameters, you are prompted for the required data.

Step 5 In another terminal window, check the installation process:

```
# tail -f /tmp/primecentral_install.log
```

Step 6 After the installation succeeds, start Prime Central:

```
# /usr/local/bin/pc.sh start
```

Step 7 Verify that Prime Central is running correctly; then, stop it:

```
# /usr/local/bin/pc.sh stop
```

Step 8 Remove the virtual IP addresses:

```
# ip addr del 192.168.1.130 dev eth0
```

Step 9 Unmount the shared partitions:

```
# umount /opt/pc
```

Setting Up the Prime Central Cluster Service

To set up and manage a cluster, you can use the CLI or the GUI. This section explains how to use the CLI. To use the GUI, see the *Red Hat Enterprise Linux 5 Cluster Administration Guide*, sections 3 and 4.

To set up the Prime Central cluster service, complete the following steps in parallel on both nodes, except where noted:

Step 1 Modify the `/etc/cluster/cluster.conf` file by setting unique values for the parameters listed in [Table 1-5](#).

Step 2 Copy the edited `cluster.conf` file to the `/etc/cluster/` directory.

Whenever you modify the `cluster.conf` file, increment the `config_version` value so the `cluster.conf` file propagates correctly to the nodes. To propagate the `cluster.conf` file manually, you must stop the cluster, copy the file, and then start the cluster.

Step 3 Start the cluster services:

```
# service cman start
# service rgmanager start
```

Enter each command on one node and then immediately enter the same command on the other node.

For example, when `cman` starts on a node, it waits for the other node to start `cman`. If the other node takes too long to start `cman`, `cman` times out on the first node.

Step 4 (For the RHCS luci GUI only) Using the username `admin`, start the RHCS ricci service:

```
# service ricci start
```

Step 5 (For the RHCS luci GUI only) On the first node only, start the RHCS luci services:

```
# luci_admin init
# service luci start
```

Step 6 To test failover, relocate the service to another node:

```
# clusvcaadm -r vmppcservice -m prime-ha-node2.cisco.com
```

Step 7 After the Prime Central service is running in an HA cluster, you cannot restart its components (such as the portal, integration layer, and database) without first freezing the cluster. After you restart the component, you can unfreeze the cluster.

For example, attaching or detaching an application to or from Prime Central requires an integration layer restart. On the active node, freeze the HA cluster, restart the integration layer, and unfreeze the cluster:

```
# clusvcadm -Z Prime-Central-service-name
# su - primeusr
# itgctl stop
# itgctl start
# exit
# clusvcadm -U Prime-Central-service-name
```

Modifying Parameters in the cluster.conf File

The following table lists the parameters in the `/etc/cluster/cluster.conf` file for which you must set unique values.

Table 1-5 Parameters to Modify in the `cluster.conf` File

Parameter	Default Value	Notes
Cluster name	bmlcluster	—
Multicast address	224.0.0.251	The multicast address must be unique per subnet and must be working before you start your cluster. For a tool to verify that your multicast address is correct, see Troubleshooting, page 1-35 .
Service IP address	192.168.1.130	—
First node name	prime-ha-node1.cisco.com	—
Second node name	prime-ha-node2.cisco.com	—
Shared partition path	/dev/mapper/mpath2p1	—

Checking the Cluster Services

On both nodes, check the status of the cluster:

```
# clustat
```

The output is similar to the following:

```
Cluster Status for bmlcluster @ Thu Jun 27 23:58:35 2013
```

```
Member Status: Quorate
```

```
Member Name          ID    Status
-----
prime-ha-node1.cisco.com 1    Online, Local, rgmanager
prime-ha-node2.cisco.com 2    Online, rgmanager

Service Name          Owner (Last)          State
-----
service:vmppcservice prime-ha-node2.cisco.com started
```

Next Steps

Complete the following steps on both nodes, except where noted:

-
- Step 1** Configure `cman`, `rgmanager`, and `ricci` to start automatically upon bootup:
- ```
chkconfig cman on
chkconfig rgmanager on
chkconfig ricci on
```
- Step 2** (On the first node only) Configure `luci` to start automatically upon bootup:
- ```
# chkconfig luci on
```
- Step 3** Verify that the required ports are open. For a list of ports that Prime Central requires, see “Prime Central Protocols and Ports” in the *Cisco Prime Central 1.4.1 Quick Start Guide*.
- Step 4** Enable the firewall:
- ```
service iptables start
chkconfig iptables on
service ip6tables start
chkconfig ip6tables on
```
- Step 5** Disable Security-Enhanced Linux (SELinux):
- ```
# vi /etc/selinux/config
SELINUX=disabled
```
-

Installing Prime Central Fault Management in a Local Redundancy HA Configuration

Installing the Prime Central Fault Management component in a dual-node, RHCS HA configuration is a three-part process:

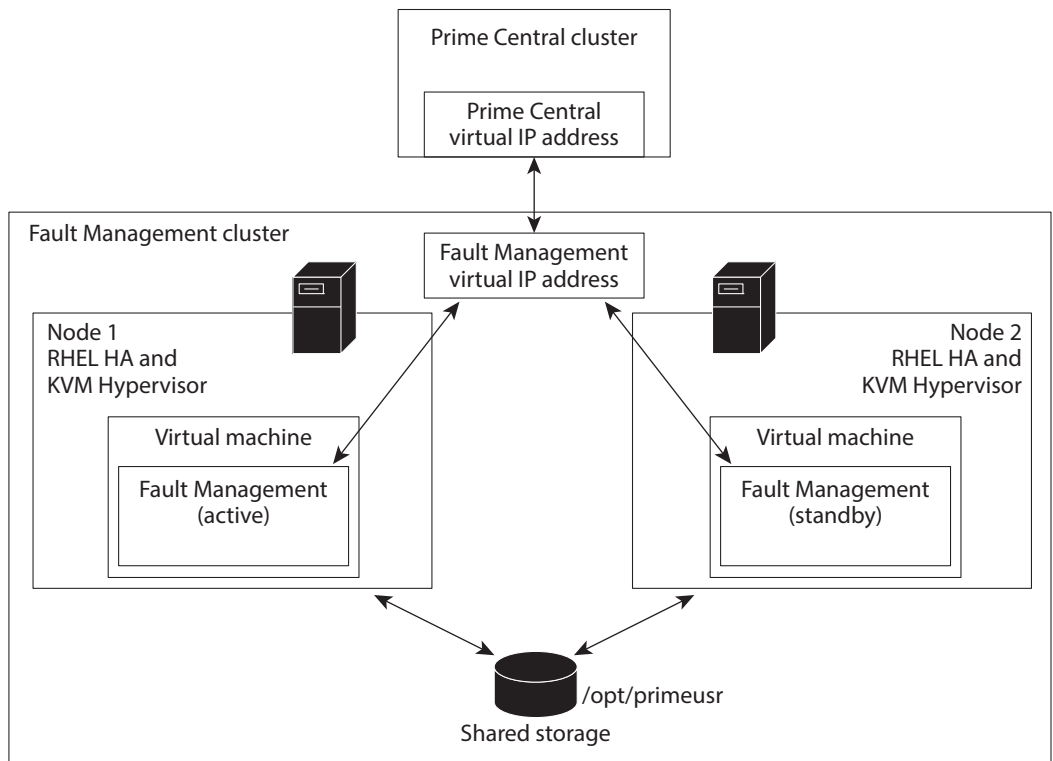
1. Install Red Hat Enterprise Linux 6.5 (RHEL 6.5) with HA and kernel-based virtual machine (KVM) packages on each node.
2. Create a single virtual machine installed with RHEL 5.8/6.5 and running the Prime Central Fault Management component.
3. Use multipath shared storage that contains the virtual machine image.

The examples provided use the following hostnames and IP addresses; yours will be different:

- Node 1—`fm-ha-node1.cisco.com` (192.168.1.150)
- Node 2—`fm-ha-node2.cisco.com` (192.168.1.160)
- Virtual IP address—`fm-service.cisco.com` (192.168.1.170)
- Gateway—192.168.1.1
- DNS—192.168.1.2

Figure 1-7 shows an example of a Fault Management cluster in an HA configuration.

Figure 1-7 Fault Management Cluster in an HA Configuration



406000

Before You Begin

- Verify that your system meets all the hardware and software requirements in “Installation Requirements” in the [Cisco Prime Central 1.4.1 Quick Start Guide](#).
- If you changed the default installation folder (`/opt/primeusr/faultmgmt`), make the equivalent changes in the following files (look for the section titled “Require manual definition” in each file):
 - `/usr/local/bin/fm.sh`
 - `/images/fm_status.sh`

Installing RHEL 6.5

To install RHEL 6.5, complete the following steps in parallel on both nodes, except where noted:

-
- Step 1** Configure specialized storage devices, high availability, and virtualization. See the Red Hat documentation for instructions.
- Step 2** Verify that the following options are checked:
- Virtualization: Virtualization Tools
 - High Availability: High Availability
 - Desktops: General Purpose Desktop

- Desktops: X Window System

Step 3 Create local directories named /rhel and /cdrom-6.5.

Step 4 Copy the .iso file that was used for the node installation to the /rhel directory.

Step 5 Mount the /rhel .iso file to /cdrom-6.5:

```
# cd /rhel
# mount -t iso9660 -o loop /rhel/rhel-server-6.5-x86_64-dvd.iso /cdrom-6.5
```



Note To permanently mount the drive, update the /etc/fstab file. See http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/4/html/Introduction_To_System_Administration/s2-storage-mount-fstab.html.

Step 6 Create a file named /etc/yum.repos.d/local.repo. Use UNIX format and be sure there are no spaces before lines.

Step 7 Save the newly created file in local.repo, as follows:

```
[local]
name=Red Hat Enterprise Linux $releasever - $basearch - Local
baseurl=file:///cdrom-6.5/Server
enabled=1
gpgcheck=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
[HighAvailability]
name=Red Hat Enterprise Linux $releasever - $basearch - HighAvailability
baseurl=file:///cdrom-6.5/HighAvailability
enabled=1
gpgcheck=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
[ResilientStorage]
name=Red Hat Enterprise Linux $releasever - $basearch - ResilientStorage
baseurl=file:///cdrom-6.5/ResilientStorage
enabled=1
gpgcheck=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
```

Step 8 (Optional) If you forget the HA package and want to install it later, enter:

```
# yum groupinstall "High Availability"
```

Step 9 (Optional) If you forget the desktop and want to install it later, enter:

```
# yum groupinstall "X Window System" Desktop
# vi /etc/inittab
```

Then, change id:3:initdefault: to **id:5:initdefault:** and reboot the server.

Step 10 Temporarily disable the firewall and SELinux to enable initial testing of the cluster:

a. To disable the firewall, enter:

```
# service iptables save
# service iptables stop
# chkconfig iptables off
# service ip6tables save
# service ip6tables stop
# chkconfig ip6tables off
```

- b. To disable SELinux, enter:

```
# vi /etc/selinux/config
change
SELINUX=enforcing
to
SELINUX=disabled
```

- Step 11** Keep the nodes synchronized:

```
# echo server tick.redhat.com$'\n'restrict tick.redhat.com mask 255.255.255.255 nomodify
notrap noquery >> /etc/ntp.conf
# chkconfig ntpd on
# service ntpd start
```

- Step 12** Switch network daemons:

```
# service NetworkManager stop
# chkconfig NetworkManager off
# yum remove NetworkManager
# chkconfig network on
```

- Step 13** Edit the /etc/hosts file to add the node information; for example:

```
192.168.1.150 prime-ha-node1.cisco.com prime-ha-node1
192.168.1.160 prime-ha-node2.cisco.com prime-ha-node2
```

- Step 14** Generate an SSH key for the root user:

```
# chmod 755 ~
# ssh-keygen -t rsa -N "" -b 2047 -f ~/.ssh/id_rsa
# chmod 600 ~/.ssh/id_rsa
```

- Step 15** (On the first node only) Share the node's public key with the other node so that dynamically creating a secure shell between the nodes does not prompt for a password:

```
# rsync -av ~/.ssh/id_rsa.pub ~/.ssh/authorized_keys
# ssh root@node2 "cat ~/.ssh/id_rsa.pub" >> ~/.ssh/authorized_keys
# rsync -av ~/.ssh/authorized_keys root@node2:/root/.ssh/
```

- Step 16** Verify that the .ssh directory has 700 permission and the .ssh/id_rsa file has 600 permission:

```
# chmod 700 ~/.ssh
# chmod 600 ~/.ssh/id_rsa
```

- Step 17** Verify that your SSH is working without an authentication or password prompt:



Caution

The Fault Management service will not start if SSH prompts for authentication or a password. Be sure to complete all of the following substeps.

- a. On node prime-ha-node1.cisco.com, enter:

```
# ssh root@prime-ha-node2.cisco.com
# exit
# ssh root@prime-ha-node2
# exit
# ssh root@192.168.1.150
# exit
```

- b. On node `prime-ha-node2.cisco.com`, enter:

```
# ssh root@prime-ha-node1.cisco.com
# exit
# ssh root@prime-ha-node1
# exit
# ssh root@192.168.1.160
# exit
```

- c. If you are prompted for a password, check the permissions of all folders and files that you modified in the preceding steps.
- d. If you are prompted to continue connecting, enter **yes**. (The prompt should appear only the first time you use SSH to connect to the node.)

Configuring Multipath

To configure multipath, complete the following steps in parallel on both nodes, except where noted:



Note

These steps set up a *nonclustered* drive. If you want to do a live migration of your virtual machines, you must set up a *clustered* drive such as the Clustered Logical Volume Manager (CLVM).

- Step 1** Install multipath:

```
# yum install device-mapper-multipath
```

- Step 2** Configure and start the services:

```
# mpathconf --enable --user_friendly_names y --find_multipaths y
# modprobe dm_multipath
# service multipathd start
# chkconfig multipathd on
```

- Step 3** Check for available disks. The names of the multipath disks must be identical on both nodes:

```
cd /dev/mapper/
ls -la
total 0
drwxr-xr-x.  2 root root   160 Jul 12 19:28 .
drwxr-xr-x. 18 root root  4160 Jul 12 19:28 ..
crw-rw----.  1 root root  10, 58 Jul 12 19:09 control
lrwxrwxrwx.  1 root root    7 Jul 12 19:28 mpathc -> ../dm-3
lrwxrwxrwx.  1 root root    7 Jul 12 19:09 vg_primehanode2-lv_home -> ../dm-2
lrwxrwxrwx.  1 root root    7 Jul 12 19:09 vg_primehanode2-lv_root -> ../dm-0
lrwxrwxrwx.  1 root root    7 Jul 12 19:09 vg_primehanode2-lv_swap -> ../dm-1
```

In the output, note *mpathc*, which is the multipath virtual device or disk that you will use later as shared storage.

Adding Shared Partitions

To add shared partitions, complete the following steps in parallel on both nodes, except where noted:



Note The examples provided use device mapping names such as `mpathc` and `mpathcpl`; yours will be different.

Step 1 (On the first node only) Create a 100-GB, shared partition:

```
# cd /dev/mapper
# fdisk mpathc
Command (m for help): p
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-19581, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-19581, default 19581): +100GB
Command (m for help): w
```

Step 2 Reboot both nodes.

Step 3 Check for new partitions:

```
# cd /dev/mapper/
# ls -la
total 0
drwxr-xr-x  2 root root    180 May  4 19:49 .
drwxr-xr-x 13 root root   4120 May  4 19:49 ..
crw-----  1 root root    10, 63 May  4 19:49 control
brw-rw----  1 root disk 253,  2 May  4 19:49 mpathc
brw-rw----  1 root disk 253,  3 May  4 19:49 mpathcpl
brw-rw----  1 root disk 253,  0 May  4 19:49 VolGroup00-LogVol100
brw-rw----  1 root disk 253,  1 May  4 19:49 VolGroup00-LogVol01
```

Step 4 Create target locations on both nodes:

```
# mkdir /images
```

Step 5 Check if the new partition is mapped to another server:

```
# mount /dev/mapper/mpathcpl /images
```

If the mount fails due to an invalid file type, the partition is not a link to an existing partition; skip to [Step 6](#).

Otherwise, run a directory listing of `/images`. If the listing contains data from an existing partition, *do not* reformat this partition. Instead, leave this partition as is and return to [Step 1](#) to create another partition.

Step 6 (On the first node only) Format the new shared partition:

```
# mkfs.ext4 /dev/mapper/mpathcpl
```

Step 7 Verify that both nodes can mount and unmount the shared storage:

- a. On the first node, mount the shared storage and save a file that contains only the value *1* to the shared storage. The test.txt file should exist in the list of contents of /images:

```
# mount /dev/mapper/mpathcp1 /images
# vi /images/test.txt
1
:wq
# ls -la /images
# umount /images
```

- b. On the second node, mount the shared storage and verify that the test.txt file exists and contains the value *1*:

```
# mount /dev/mapper/mpathcp1 /images
# vi /images/test.txt
:q
# umount /images
```

If you cannot mount or unmount the shared storage, or if the test.txt file does not exist when you mount it to the second node, your multipath is not set up correctly.

Setting Up the Prime Central Fault Management Virtual Machine

Step 1 Mount the newly created partition on the first node:

```
# mount /dev/mapper/mpathcp1 /images
```

Step 2 (On the first node only) Add a new storage pool:

- Run **vncserver** and use the VNC client to access the node.
- Launch the virt-manager.
- Click **Edit Connection Details**.
- Click the **Storage** tab.
- Click the + button to add a new storage pool.
- In the Add a New Storage Pool: Step 1 of 2 window, enter **fm_images** as the storage pool name, choose **fs: Pre-Formatted Block Device** as the type, and click **Forward**.
- In the Add a New Storage Pool: Step 2 of 2 window, verify that the settings are as follows; then, click **Finish**:
 - Target Path: **/images**
 - Format: **auto**
 - Source Path: **/dev/mapper/mpathcp1**

Step 3 (On the first node only) Add a new storage volume:

- In the virt-manager, click the **Storage** tab.



Caution

Do not check the Autostart: On Boot check box.

- b. Click the **New Volume** button.
- c. In the Add a Storage Volume window, enter the following values; then, click **Finish**:
 - Name: **fm_vm** (.img is appended)
 - Format: **raw**
 - Max Capacity (MB): Use all available storage space from the pool.
 - Allocation (MB): **0**

Step 4 Create a virtual network:

- a. On each node, add a bridge to the host, to enable the virtual machines to use the same physical network as the nodes:

```
# cd /etc/sysconfig/network-scripts/
# vi ifcfg-br0
```

Add these lines and save:

```
DEVICE=br0
TYPE=Bridge
BOOTPROTO=static
ONBOOT=yes
DELAY=0
IPADDR=192.168.1.150
GATEWAY=192.168.1.1
DNS1=192.168.1.2
```



Note The IPADDR has the same value as the node to which you are adding this file. This example is for node 1.

```
# vi ifcfg-eth0
```

Add this line and save:

```
BRIDGE=br0

# service network restart
# brctl show
bridge name      bridge id                STP enabled  interfaces
br0              8000.0025b500005b       no           eth0
                vnet0
virbr0          8000.5254003af3e9       yes          virbr0-nic
```

- b. On each node, update the `/etc/sysctl.conf` file to allow forwarding to the virtual machines:

```
# vi /etc/sysctl.conf
net.ipv4.ip_forward = 1
```

Step 5 Create a new virtual machine:

- a. Copy the RHEL 5.8/6.5 .iso file to the `/rhel` directory.
- b. In the virt-manager window, click the **create a new virtual machine** button.
- c. In the Step 1 of 5 window, enter **fm_vm** as the virtual machine name, click **Local install media**, and click **Forward**.
- d. In the Step 2 of 5 window, click **Use ISO image** and specify the location of the RHEL 5.8/6.5 .iso image. Verify that the OS type is **Linux** and the version is **Red Hat Enterprise Linux 5.4 or later**. Then, click **Forward**.

- e. In the Step 3 of 5 window, enter the amount of RAM and CPUs to use for the virtual machine. For recommendations, see “Installation Requirements” in the *Cisco Prime Central 1.4.1 Quick Start Guide*. Then, click **Forward**.
- f. In the Step 4 of 5 window, check **Enable storage for this virtual machine**. Click **Select managed or other existing storage** and browse to /images/fm_vm.img (which you created in [Step 3c](#)). Then, click **Forward**.
- g. In the Step 5 of 5 window, verify that the settings are as follows; then, click **Finish**:
 - Advanced options: Host device eth0 (Bridge 'br0')
 - Virt Type: kvm
 - Architecture: x86_64

Step 6 Install RHEL 5.8/6.5 on the new virtual machine. The OS library, kernel-headers, is required; choose the **Software Development** option when installing RHEL to ensure that the kernel-headers library is installed.

Step 7 Temporarily disable the firewall and SELinux to enable initial testing of the cluster:

- a. To disable the firewall, enter:

```
# service iptables save
# service iptables stop
# chkconfig iptables off
# service ip6tables save
# service ip6tables stop
# chkconfig ip6tables off
```

- b. To disable SELinux, enter:

```
# vi /etc/selinux/config
change
SELINUX=enforcing
to
SELINUX=disabled
```

Step 8 Update the /etc/hosts file on the virtual machine:

```
# IP-address FQDN hostname
```

For example:

```
192.168.1.170 fm-service.cisco.com fm-service
```

Step 9 From the virtual machine, ping both nodes. If the ping fails, add both nodes to the virtual machine’s /etc/hosts file. For example:

```
192.168.1.150 prime-ha-node1.cisco.com prime-ha-node1
192.168.1.160 prime-ha-node2.cisco.com prime-ha-node2
```

Step 10 Save the /etc/hosts file; then, run the following tests:

```
# hostname -a
fm-service
# hostname -f
fm-service.cisco.com
# hostname -i
192.168.1.170
# ipcalc -h 192.168.1.170
HOSTNAME=fm-service.cisco.com
```

If any of the tests return incorrect results, check the `/etc/hosts` file for typos. Also check the `/etc/sysconfig/network` file and verify that the `HOSTNAME` entry contains your server's FQDN (`fm-service.cisco.com` in this example).

Step 11 Generate an SSH key for the virtual machine's root user and share it with both nodes:

```
# chmod 755 ~
# ssh-keygen -t rsa -N "" -b 2047 -f ~/.ssh/id_rsa
# chmod 600 ~/.ssh/id_rsa
# rsync -av ~/.ssh/id_rsa.pub ~/.ssh/authorized_keys
# ssh root@prime-ha-node1.cisco.com "cat ~/.ssh/id_rsa.pub" >> ~/.ssh/authorized_keys
# ssh root@prime-ha-node2.cisco.com "cat ~/.ssh/id_rsa.pub" >> ~/.ssh/authorized_keys
# rsync -av ~/.ssh/authorized_keys root@prime-ha-node1.cisco.com:/root/.ssh/
# rsync -av ~/.ssh/authorized_keys root@prime-ha-node2.cisco.com:/root/.ssh/
```

Step 12 On the virtual machine, verify that the `.ssh` directory has 700 permission and the `.ssh/id_rsa` file has 600 permission:

```
# chmod 700 ~/.ssh
# chmod 600 ~/.ssh/id_rsa
```

Step 13 Verify that your SSH is working without an authentication or password prompt:



Caution

The Fault Management service will not start if SSH prompts for authentication or a password. Be sure to complete all of the following substeps.

a. On node `prime-ha-node1.cisco.com`, enter:

```
# ssh root@fm-service.cisco.com
# exit
# ssh root@fm-service
# exit
# ssh root@192.168.1.170
# exit
```

b. On node `prime-ha-node2.cisco.com`, enter:

```
# ssh root@fm-service.cisco.com
# exit
# ssh root@fm-service
# exit
# ssh root@192.168.1.170
# exit
```

c. If you are prompted for a password, check the permissions of all folders and files that you modified in the preceding steps.

d. If you are prompted to continue connecting, enter **yes**. (The prompt should appear only the first time you use SSH to connect to the node.)

Step 14 Distribute the virtual machine:

a. Click the running virtual machine and choose **Shutdown > Save**.

b. On the first node, copy the virtual machine definitions file to the shared directory:

```
# virsh dumpxml fm_vm > /images/fm_vm.xml
# mkdir /var/images
# cp /images/fm_vm.xml /var/images/fm_vm.xml
# umount /images
```

- c. On the second node, copy the virtual machine definition to the second node:

```
# mount /dev/mapper/mpathcp1 /images
# virsh define /images/fm_vm.xml
# mkdir /var/images
# cp /images/fm_vm.xml /var/images/fm_vm.xml
# umount /images
```

Distributing the Node .tar File

Step 1 Insert the Cisco Prime Central 1.4.1 USB, navigate to the High Availability/RHCS Bare Metal Local HA/Fault Management folder, and locate the primefm_v1.4.1_ha_node.tar.gz file.

Step 2 Use SSH to connect to the first node.

Step 3 Copy the primefm_v1.4.1_ha_node.tar.gz file to the first node.

Step 4 Back up the /etc/cluster/ and /images directories.

Step 5 Distribute the file:

```
# mount /dev/mapper/mpathcp1 /images
# tar -zxf primefm_v1.4.1_ha_node.tar.gz -C / --owner root --no-same-owner
# chmod 755 /images/*.sh
```

Step 6 Edit the fm_status.sh file by changing VM_FQDN to your virtual machine's FQDN:

```
# vi /images/fm_status.sh
```

For example:

```
VM_FQDN=fm-service.cisco.com
```

Distributing the Virtual Machine .tar File

Step 1 Insert the Cisco Prime Central 1.4.1 USB, navigate to the High Availability/RHCS Bare Metal Local HA/Fault Management folder, and locate the primefm_v1.4.1_ha_vm.tar.gz file.

Step 2 Mount the shared drive. (The shared storage should still be mounted to the first node. If not, verify that the shared storage is not mounted to the other node; then, mount it to the first node.)

```
# mount /dev/mapper/mpathcp1 /images
```

Step 3 Launch the virtual machine:

```
# virsh start fm_vm
```

Step 4 Use SSH to connect to the virtual machine.

Step 5 Copy the primefm_v1.4.1_ha_vm.tar.gz file to the virtual machine.

Step 6 Back up the /root/ha-stuff/fm and /usr/local/bin directories.

Step 7 Distribute the file:

```
# cd /root/ha-stuff/fm
# tar -zxf FM1.4.1Build.tar.gz
# cp /root/ha-stuff/fm/fm_install.properties /root/ha-stuff/fm/Disk1/InstData/VM/
# chmod -R 755 /root/ha-stuff/fm/Disk*
# chmod 755 /usr/local/bin/*.sh
```

Installing Prime Central Fault Management on the Virtual Machine

You can use the GUI to install Fault Management on the virtual machine, or you can install it silently. The following procedures explain both the GUI-based and silent installation options; choose your preferred option.

GUI-Based Installation

Step 1 Use SSH to connect to the virtual machine.

Step 2 Enter:

```
# cd /root/ha-stuff/fm/Disk1/InstData/VM
```

Step 3 Install Prime Central Fault Management:

```
# ./primefm_v1.4.1.bin
```

At the installer prompts, enter the information described in the [Cisco Prime Central 1.4.1 Quick Start Guide](#), section “Installing Prime Central Fault Management on the Server.”

Step 4 In another terminal window, check the installation process:

```
# tail -f /opt/primeusr/faultmgmt/install/log/PrimeFM-*.log
```

Step 5 After the installation succeeds, use SSH to connect to the Prime Central HA active server and enter:

```
# clusvcadm -Z Prime-Central-service-name
# su - primeusr
# itgctl stop
# itgctl start
# clusvcadm -U Prime-Central-service-name
```

Step 6 To test the Prime Central Fault Management installation, open a browser, log in to the Prime Central portal, and verify that the Prime Central Fault Management component is running.

Silent Installation

Step 1 Use SSH to connect to the virtual machine.

Step 2 Enter:

```
# cd /root/ha-stuff/fm/Disk1/InstData/VM
```

Step 3 Edit the fm_install.properties file to match your setup.

- Step 4** Install Prime Central Fault Management:
- ```
./primefm_v1.4.1.bin -i silent -f fm_install.properties
```
- Step 5** In another terminal window, check the installation process:
- ```
# tail -f /opt/primeusr/faultmgmt/install/log/PrimeFM-*.log
```
- Step 6** The silent installation does not report errors. To see if any errors occurred, check the log files—starting with primefm.log—in the /opt/primeusr/faultmgmt/install/log folder.
- Step 7** After the installation succeeds, use SSH to connect to the Prime Central HA active server and enter:
- ```
clusvcadm -Z Prime-Central-service-name
su - primeusr
itgctl stop
itgctl start
clusvcadm -U Prime-Central-service-name
```
- Step 8** To test the Prime Central Fault Management installation, open a browser, log in to the Prime Central portal, and verify that the Prime Central Fault Management component is running.
- Step 9** Remove the fm\_install.properties file, which contains your server’s passwords used during the silent installation.
- 

## Setting Up the Fault Management Cluster Service

---

- Step 1** Run the 6.5 cluster workaround on both nodes (see <https://access.redhat.com/knowledge/solutions/67583>):
- ```
# mkdir /var/lib/ricci/.libvirt
# chown ricci:ricci /var/lib/ricci/.libvirt
```
- Step 2** Verify that the system contains a user named *ricci*. If the ricci user is missing, enter:
- ```
useradd ricci
```
- Step 3** On both nodes, edit and save the vm.sh agent to allow a longer start time for the Fault Management virtual machine. (If you are using different hardware, you might need to increase the timeout value.)
- Enter:
 

```
vi /usr/share/cluster/vm.sh
```
  - Locate the timeout value (in seconds):
 

```
<action name="start" timeout="300"/>
```
  - Change the timeout value to:
 

```
<action name="start" timeout="600"/>
```
- Step 4** Modify the /etc/cluster/cluster.conf file and set unique values. The multicast address must be unique per subnet and must be working before you start your cluster. For a tool to verify that your multicast address is correct, see [Troubleshooting, page 1-35](#).
- Step 5** Copy the edited cluster.conf file to the /etc/cluster/ directory on both nodes.
- If the cluster is not up and running when you change the cluster.conf file, manually copy cluster.conf to the other node; then, restart the cluster:



```
rsync -av /etc/cluster/cluster.conf root@prime-ha-node2.cisco.com:/etc/cluster/
```

**Step 6** Validate the cluster.conf file:

```
ccs_config_validate
```

**Step 7** Install luci and ricci (if they are not already installed):

```
yum install luci
yum install ricci
```

**Step 8** Start the ricci service on both nodes:

```
passwd ricci
service ricci start
```



**Note** Enter the `passwd ricci` command only once; doing so creates a password for the user `ricci`.

**Step 9** Start the cluster services on both nodes:

```
service cman start
service rgmanager start
```

**Step 10** (Only if you are using the RHCS luci GUI) On the first node only, start the luci service:

```
luci_admin init
service luci start
```

**Step 11** (Only if you are using the RHCS luci GUI) Log in to luci on the node where you started the luci service; for example:

```
https://prime-ha-node1.cisco.com:8084
```

## Checking the Cluster Services

**Step 1** Review the cluster log file in `/var/log/messages`.

**Step 2** Check the status of the cluster:

```
clustat
```

The output is similar to the following:

```
Cluster Status for bmlcluster @ Tue Jul 30 15:47:46 2013
Member Status: Quorate
```

Member Name	ID	Status
prime-ha-node1.cisco.com	1	Online, Local, rgmanager
prime-ha-node2.cisco.com	2	Online, rgmanager

Service Name	Owner (Last)	State
service:vm1	prime-ha-node1.cisco.com	started

- Step 3** Test the Prime Central Fault Management installation:
- Open a browser, log in to the Prime Central portal, and verify that the Prime Central Fault Management component is running.
  - Relocate the virtual machine:
 

```
clusvcdm -r vml -m prime-ha-node2.cisco.com
```
  - After the relocation is complete, reverify that the Fault Management component is running on the Prime Central portal.
- Step 4** After the Fault Management service is running in an HA cluster, you cannot restart its components (such as Netcool/Impact, OMNIbus, and Tivoli Common Reporting [TCR]) without first freezing the cluster. After you restart the component, you can unfreeze the cluster.
- To restart a Fault Management component:
- On the active Fault Management node, enter:
 

```
clusvcdm -z Fault-Management-service-name
```
  - Use SSH to connect to the Fault Management virtual machine and enter:
 

```
/usr/local/bin/fm.sh stop
/usr/local/bin/fm.sh start
```
  - Use SSH to connect to the active Fault Management node and enter:
 

```
clusvcdm -U Fault-Management-service-name
```
- 

## Next Steps

Complete the following steps on both nodes, except where noted:

- Step 1** Configure cman, rgmanager, and ricci to start automatically upon bootup:
- ```
# chkconfig cman on
# chkconfig rgmanager on
# chkconfig ricci on
```
- Step 2** (On the first node only) Configure luci to start automatically upon bootup:
- ```
chkconfig luci on
```
- Step 3** Verify that the required ports are open. For a list of ports that the Fault Management component requires, see “Prime Central Protocols and Ports” in the [Cisco Prime Central 1.4.1 Quick Start Guide](#).

**Step 4** (On both nodes and on the virtual machine) Enable the firewall:

```
service iptables start
chkconfig iptables on
service ip6tables start
chkconfig ip6tables on
```

**Step 5** Disable SELinux:

```
vi /etc/selinux/config
SELINUX=disabled
```

---

## Troubleshooting

The following troubleshooting steps help solve common problems in an HA configuration.

**Problem** The HA installation fails.

**Solution** Check the log files to locate the problem and take the appropriate action. Log files contain detailed information about request processing and exceptions and are your best diagnostic tool for troubleshooting. See “Troubleshooting the Installation” in the *Cisco Prime Central 1.4.1 Quick Start Guide*.

**Problem** Prime Central does not start in a clustered setup.

**Solution** Check the `/var/log/messages` files for failure to either mount the shared storage or add the virtual IP address. If the shared storage failed to mount, shut down the cluster and verify that you can manually add the shared storage to a node. (Be sure to unmount it after your test.)

If the virtual IP address was not added, verify that it is in the same subnet as the nodes and is not in use by any other computer in the network.

If you find that `/usr/local/bin/pc.sh` start failed, check `/usr/local/bin/pc.log` and `/usr/local/bin/pc-start.log`, which will tell you if the database or other Prime Central components failed to start. Then, to determine which component failed to start:

1. Stop the `luci`, `ricci`, `rgmanager`, and `cman` services on both nodes to shut down the cluster.
2. On the node where you originally installed Prime Central:

- a. Mount the shared storage.
- b. Add the virtual IP address.
- c. Verify that all services have stopped:

```
/usr/local/bin/pc.sh stop
```

- d. Enter:

```
su - primeusr
emdbctl -start
itgctl start
portalctl start
```

- e. Check the output from each of the preceding commands to locate the problem.

**Problem** You receive the error “<err> 'fsck -p /dev/mapper/mpath2p1' failed, error=4; check /tmp/fs-vmpecfs.fsck.log.mq4986 for errors.”

**Solution** Enter the following command and reboot when it is finished running:

```
fsck -f /dev/mapper/mpath2p1
```

**Problem** You receive the error “Timeout exceeded while waiting for ‘/images/fm\_status.sh’” in /var/log/messages.

**Solution** Verify that you can use SSH to connect to each node and virtual machine without an authentication or password prompt. If SSH prompts for authentication or a password, the Prime Central and Fault Management services cannot start.

**Problem** Your environment uses the wrong fencing device.

**Solution** The examples in this guide use fence\_manual and fence\_virsh, which are test fencing devices and cannot be used for production. For information about which fencing device to use in your environment, see the *Red Hat Enterprise Linux 6 Cluster Administration: Configuring and Managing the High Availability Add-On*.

**Problem** The cman and rgmanager services do not start.

**Solution** Check the log files in /var/log/messages and /var/log/cluster. Use the following tool to verify that your multicast address is correct:  
<http://juliandyke.wordpress.com/2010/12/03/testing-multicasting-for-oracle-11-2-0-2-grid-infrastructure/>.

**Problem** Cannot perform a live migration.

**Solution** To support live migration of the virtual machines, confirm that the shared storage is set up as a clustered file system, such as Global File System (GFS) or CLVM.

**Problem** Cannot stop the cluster.

**Solution** Use luci or the command line to shut down your cluster:

- luci—Select the cluster; then, from the drop-down list, choose **Stop this cluster**.
- Command line—Alternating between the two nodes, shut down the services in the reverse order in which you started them. For example, enter the **stop** command for rgmanager on node1; then, enter it on node2. Enter the **stop** command for cman on node1; then, enter it on node2.

```
service luci stop
service ricci stop
service rgmanager stop
service cman stop
```

**Problem** When trying to unmount the shared storage, a “device is busy” message is returned.

**Solution** Verify that all cluster services have stopped and that you have closed all terminal sessions that are accessing the shared storage location. To determine which user is accessing the shared storage, enter:

```
fuser -m -v shared-storage
```

For example:

```
fuser -m -v /opt/pc
```

**Problem** You do not know if the node can support virtualization.

**Solution** Enter:

```
egrep '^flags.*(vmx|svm)' /proc/cpuinfo
```

If the command returns no output, the node does not support virtualization.

If the command output contains `vmx` or `svm` flags, the node supports virtualization. For example:

```
flags : fpu vme de pse tsc msr pae mce cx8 apic mtrr pge mca cmov pat pse36
clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm constant_tsc
arch_perfmon pebs bts rep_good xtopology nonstop_tsc aperfmperf pni pclmulqdq dtes64
monitor ds_cpl vmx smx est tm2 sse3 cx16 xtpr pdcm dca sse4_1 sse4_2 popcnt aes lahf_lm
ida arat dts tpr_shadow vnmi flexpriority ept vpid
```

**Problem** You receive the error “operation failed: domain 'fm\_vm' already exists with uuid...”

**Solution** An `fm_vm.xml` file might already exist on the second node due to a previous attempt to define the file. Do the following:

1. Verify that `/images` is unmounted from the first node.
2. On the second node (that is, the node on which you did *not* create the virtual machine), enter:

```
mv /etc/libvirt/qemu/fm_vm.xml /tmp
mount /dev/mapper/mpathcpl /images
virsh define /images/fm_vm.xml
umount /images
```

**Problem** Cannot test the `cluster.conf` file.

**Solution** Use `rg_test` commands. For example:

- To display the resource rules that `rg_test` understands, enter:

```
rg_test rules
```

- To test a configuration, enter:

```
rg_test test /etc/cluster/cluster.conf
```

- To display the start ordering of a service, enter:

```
rg_test noop /etc/cluster/cluster.conf start service service-name
```

- To display the stop ordering of a service, enter:

```
rg_test noop /etc/cluster/cluster.conf stop service service-name
```

**Problem** When you reboot one or both nodes, the node is fenced before it can join the cluster.

**Solution** To start up, the node might require an additional fencing delay. Edit your `cluster.conf` file by increasing the value of the `post_join_delay` attribute:

```
<fence_daemon clean_start="0" post_fail_delay="0" post_join_delay="30"/>
```

**Problem** After you relocate the Prime Central service, the integration layer is shown in the Prime Central Suite Monitoring portlet > Applications tab, but its state is Down.

**Solution** On servers where the hardware requirements are at or below the minimum for Prime Central high availability, the integration layer requires more time to start up. Do the following:

1. On the active node where Prime Central is running, locate the `/opt/pc/primecentral/esb/etc/com.cisco.prime.esb.jms.cfg` file.
2. Edit the file by increasing the `waitForStart` attribute for the `jmsvm.internalBrokerURL` property. (If the line is commented, uncomment it.)

The default `waitForStart` value is 10,000 milliseconds; increase it depending on the slowness of your server. For example, to increase the `waitForStart` value to 30 seconds, enter:

```
jmsvm.internalBrokerURL=vm://internalBroker?broker.persistent=false&jms.prefetchPolicy.queuePrefetch=1&create=false&waitForStart=30000
```

**Problem** The Prime Central portal does not look correct.

**Solution** The cluster manager might have relocated the server. Clear your browser cache and refresh your screen; then, log back into the Prime Central portal.

**Problem** You need to restart a Prime Central or Fault Management component in an HA environment.

**Solution** Prime Central contains components such as the portal, integration layer, and database. Fault Management contains components such as Netcool/Impact, OMNIBus, and TCR. If you need to perform maintenance on a specific component, you must freeze the HA cluster before you can stop the component. After you restart the component, you can unfreeze the cluster.

- To freeze the cluster, enter:  
`clusvcadm -Z service-name`
- To unfreeze the cluster, enter:  
`clusvcadm -U service-name`

**Problem** After adding multipath, you cannot see the multipath names when listing the `/dev/mapper` directory.

**Solution** Do the following:

1. Enter:  
`# vi /etc/multipath.conf`
2. Change the `find_multipaths` value to **no**.
3. Enter:  
`# wq`  
`# service multipathd reload`

You should now see the multipath names.

**Problem** You receive the following error while mounting the storage:

```
[root@prime-central-linux4 ~]# mount /dev/mapper/mpathbp2 /images/
mount: wrong fs type, bad option, bad superblock on /dev/mapper/mpathbp2,
 missing codepage or helper program, or other error
 In some cases useful info is found in syslog - try
 dmesg | tail or so
```

**Solution** Enter:

```
fsck.ext3 <storage-path>
fsck.ext3 /dev/mapper/mpathbp2
```

**Problem** The `fmctl status` command shows that Fault Management started in KVM, but hangs in “starting” status in the cluster.

**Solution** Check the SSH password-less connection between the two nodes and KVM.

## Upgrading to Prime Central 1.4.1 in a Local Redundancy HA Configuration

- 
- Step 1** Shut down the cluster.
- Step 2** Insert the Cisco Prime Central 1.4.1 USB, navigate to the High Availability/RHCS Bare Metal Local HA/Prime Central folder, and locate the `primecentral_v1.4.1_ha_vm.tar.gz` file.
- Step 3** Use SSH to connect to the first node.
- Step 4** Copy the `primecentral_v1.4.1_ha_vm.tar.gz` file to the first node.
- Step 5** Back up the following directories on both nodes:
- `/root/ha-stuff/pc`
  - `/usr/local/bin`
- Step 6** Distribute the file:
- ```
# tar -zxvf primecentral_v1.4.1_ha_vm.tar.gz -C / --owner root --no-same-owner
```
- Step 7** Navigate to the Base Application folder and copy `primecentral_v1.4.1.bin` and all available `.zip` files to the `/root/ha-stuff/pc` directory:
- ```
chmod 755 /usr/local/bin/*
chmod 755 /root/ha-stuff/pc/*
```
- Step 8** On the first node only, do the following:
- a. Mount the shared partitions:
 

```
mount /dev/mapper/mpath2p1 /opt/pc
```
  - b. Add a virtual IP cluster service address for the Prime Central service:
 

```
ip addr add 192.168.1.130 dev eth0
```

- c. Update the `install.properties` file and verify that all required properties have values. Review the comments at the top of the `install.properties` file for details.



**Note** To install Prime Central silently, you must edit the `/root/ha-stuff/pc/install.properties` file. See “Sample `install.properties` Files” in the [Cisco Prime Central 1.4.1 Quick Start Guide](#).

- d. Install Prime Central:

```
cd /root/ha-stuff/pc
./PrimeCentral.sh 192.168.1.130 node's-root-password second-node-IP-address
```



**Note** You run the `PrimeCentral.sh` script by adding the preceding command-line parameters. If you do not add the command-line parameters, you are prompted for the required data.

- e. In another terminal window, check the upgrade process:

```
tail -f /tmp/primecentral_install.log
```

- f. After the upgrade succeeds, start Prime Central:

```
/usr/local/bin/pc.sh start
```

- g. Verify that Prime Central is running correctly; then, stop it:

```
/usr/local/bin/pc.sh stop
```

- h. Remove the virtual IP addresses:

```
ip addr del 192.168.1.130 dev eth0
```

- i. Unmount the shared partitions:

```
umount /opt/pc
```

When the upgrade completes, the log files are available in `~/upgrade/1.4.0.0-1.4.1.0/upgrade.log`.

## Upgrading to Prime Central Fault Management 1.4.1 in a Local Redundancy HA Configuration

- Step 1** Shut down the cluster.
- Step 2** Insert the Cisco Prime Central 1.4.1 USB, navigate to the High Availability/RHCS Bare Metal Local HA/Fault Management folder, and locate the `primefm_v1.4.1_ha_node.tar.gz` file.
- Step 3** Use SSH to connect to the first node.
- Step 4** Copy the `primefm_v1.4.1_ha_node.tar.gz` file to the first node.
- Step 5** Back up the `/etc/cluster/` and `/images` directories.
- Step 6** Distribute the file:

```
mount /dev/mapper/mpathcp1 /images
tar -zxf primefm_v1.4.1_ha_node.tar.gz -C / --owner root --no-same-owner
```



**Step 7** Navigate to the top-level Fault Management folder and copy FM1.4.1Build.tar.gz to the /root/ha-stuff/fm directory:

```
cd /root/ha-stuff/fm
tar -zxf FM1.4.1Build.tar.gz
chmod -R 755 /root/ha-stuff/fm/Disk*
chmod 755 /usr/local/bin/*.sh
chmod 755 /images/*.sh
```

**Step 8** Edit the fm\_status.sh file by changing VM\_FQDN to your virtual machine's FQDN:

```
vi /images/fm_status.sh
```

For example:

```
VM_FQDN=fm-service.cisco.com
```

**Step 9** Return to the High Availability/RHCS Bare Metal Local HA/Fault Management folder and locate the primefm\_v1.4.1\_ha\_vm.tar.gz file.

**Step 10** Mount the shared drive. (The shared storage should still be mounted to the first node. If not, verify that the shared storage is not mounted to the other node; then, mount it to the first node.)

```
mount /dev/mapper/mpathcp1 /images
```

**Step 11** Launch the virtual machine:

```
virsh start fm_vm
```

**Step 12** Use SSH to connect to the virtual machine.

**Step 13** Copy the primefm\_v1.4.1\_ha\_vm.tar.gz file to the virtual machine.

**Step 14** Back up the /root/ha-stuff/fm and /usr/local/bin directories.

**Step 15** Distribute the file:

```
tar -zxf primefm_v1.4.1_ha_vm.tar.gz -C / --owner root --no-same-owner
cd /root/ha-stuff/fm
tar -zxf FM1.4.1Build.tar.gz
cp /root/ha-stuff/fm/fm_install.properties /root/ha-stuff/fm/Disk1/InstData/VM/
chmod -R 755 /root/ha-stuff/fm/Disk*
chmod 755 /usr/local/bin/*.sh
```

**Step 16** Enter:

```
cd /root/ha-stuff/fm/Disk1/InstData/VM
```

**Step 17** Edit the fm\_install.properties file to match your setup.

**Step 18** Enter:

```
./primefm_v1.4.1.bin -i silent -f fm_install.properties
```

**Step 19** In another terminal window, check the upgrade process:

```
tail -f /opt/primeusr/faultmgmt/install/log/PrimeFM-*.log
```

**Step 20** To see if any errors occurred, check the log files—starting with primefm.log—in the /opt/primeusr/faultmgmt/upgrade/logs folder.

**Step 21** After the upgrade succeeds, use SSH to connect to the Prime Central HA active server and enter:

```
clusvadm -Z Prime-Central-service-name
su - primeusr
itgctl stop
itgctl start
clusvcadm -U Prime-Central-service-name
```

- Step 22** To test the Prime Central Fault Management upgrade, open a browser, log in to the Prime Central portal, and verify that the Prime Central Fault Management component is running.
- Step 23** Remove the `fm_install.properties` file, which contains your server's passwords.

## Uninstalling Prime Central Fault Management



### Note

If you are also uninstalling Prime Central, you must uninstall the Fault Management component first.

- Step 1** From the Prime Central portal, choose **Administration > System > Suite Monitoring > Applications** tab, and remove Fault Management.
- Step 2** Use SSH to connect to the Prime Central active node and do the following:
- Freeze the Prime Central cluster:
 

```
clusvcadm -Z service-name
```
  - Restart the integration layer:
 

```
su - primeusr
itgctl stop
itgctl start
```
  - Unfreeze the Prime Central cluster:
 

```
clusvcadm -U service-name
```
- Step 3** Use SSH to connect to the Prime Central Fault Management active node and freeze the Fault Management cluster service:
- ```
clusvcadm -Z service-name
```
- Step 4** Use SSH to connect to the Prime Central Fault Management virtual machine and do the following:
- Navigate to the `/var/adm/cisco/uninstall/Uninstall_Prime_Central_Fault_Management` directory. The `uninstall` folder contains the `installvariables.properties` file.
 - Uninstall Prime Central Fault Management:


```
./Uninstall_Prime_Central_Fault_Management -i silent
```

The uninstallation log files are available at `/var/adm/cisco/uninstall/PrimeFM-uninstall.log-time-stamp`.
- Step 5** Stop and disable the Fault Management cluster service on both nodes:
- ```
service luci stop
service ricci stop
service rgmanager stop
service cman stop
```

# Uninstalling Prime Central

---

**Step 1** Stop and disable the Prime Central cluster service on both nodes:

```
service luci stop
service ricci stop
service rgmanager stop
service cman stop
```

**Step 2** Mount and add the virtual IP address to the node that was used for the Prime Central installation:

```
ip addr add 192.168.1.130 dev eth0
mount /dev/mapper/mpath2p1 /opt/pc
```

**Step 3** Uninstall the application:

a. As the root user, log in to the Prime Central server. (If you logged in previously as a nonroot user, enter the **su - command** to become the root user.)

b. Uninstall Prime Central:

```
cd /root/ha-stuff/pc
./UninstallPrimeCentral.sh 192.168.1.120
```

In the preceding command, 192.168.1.120 is the IP address of the second node.

The uninstallation log files are available at  
`/var/adm/cisco/uninstall/UNINSTALL_LOG_time-stamp`.

**Step 4** Unmount and remove the virtual IP address:

```
ip addr del 192.168.1.130 dev eth0
umount /opt/pc
```

---

