



Installing and Maintaining Gateway Local Redundancy

The following topics provide procedures for setting up, installing, and maintaining the gateway local redundancy solution. Local redundancy is configured and monitored using the Red Hat Cluster Server (RHCS) for local redundancy. This chapter also explains how to install Prime Network Operations Reports and the Prime Network Integration Layer (PN-IL) with gateway local redundancy.



Note

Gateway high availability is supported only when the gateway software, Oracle database, and Infobright database (applicable for Operations Reports) are installed on the same server. Operations Reports are only available to customers with Operations Reports license prior to May 2018. For re-installation of Operations Reports contact a Cisco customer support representative.

- [Steps for Installing the Gateway Local Redundancy Solution, page 3-1](#)
- [Installation Requirements for Local Redundancy, page 3-4](#)
- [Preparing to Install the Local Redundancy Solution, page 3-8](#)
- [Installing the Prime Network Gateway Local Redundancy Software, page 3-15](#)
- [Verifying the Local Redundancy Setup, page 3-21](#)
- [Post-Installation Tasks for Local Redundancy, page 3-24](#)
- [Maintaining Local Redundancy, page 3-26](#)
- [Uninstalling Local Redundancy, page 3-30](#)
- [Installing and Configuring PN-IL with Local Redundancy, page 3-30](#)

Before proceeding with this chapter, make sure you have read [Local Redundancy Functional Overview, page 2-1](#).

Steps for Installing the Gateway Local Redundancy Solution

[Table 3-1](#) lists the steps you must follow to prepare for an installation, perform an installation, and verify an installation of the Prime Network gateway local redundancy solution. The table includes steps for working in a deployment that also has local redundancy. For local redundancy, the steps assume the primary database is on cluster node P1. An **x** means you must perform the step *on that server*.

**Note**

If you also have local redundancy installed, this procedure assumes the primary database is on the primary cluster server (P1).

Table 3-1 Steps for Setting Up and Installing Local Redundancy

	Task	Topic/Action Required	Primary Node P1 ¹	Standby Node P2
Step 1	Collect server details, so that you have all information handy prior to installation.	<ul style="list-style-type: none"> Prime Network Virtual IP address Oracle Virtual IP address Node 1, Node 2 Hostname and IP addresses 	x	x
Step 2	Verify that the Prime Network servers meets the prerequisites.	Installation Requirements for Local Redundancy, page 3-4	x	x
Step 3	Configure the dual-node cluster server hardware including configuring the external storage.	Configuring Hardware and External Storage for Red Hat Cluster Site, page 3-8	x	x
Step 4	Install RHEL and all recommended patches on both servers in the cluster.	Installing RHEL and Verifying the Version, page 3-9	x	x
Step 5	Install the RPMs required for Red Hat and Oracle. If you are installing Operations Reports, be sure to check this section.	Installing RPMs Required on Red Hat for Prime Network, page 3-9	x	x
Step 6	Configure disk groups, volumes, and partitions. If you are installing Operations Reports, be sure to check the required volume sizes.	Configuring Disk Group and Volumes, page 3-11	x	x
Step 7	Verify that all nodes are ready for installation by checking disk access, Linux versions, and NTP synchronization.	Verify That All Servers Are Ready for Installation, page 3-12	x	x
Step 8	Mount the external shared storage, Oracle and Prime Network mount points on the relevant directories.	Creating the Mount Points for Installation, page 3-13	x	x

Table 3-1 Steps for Setting Up and Installing Local Redundancy (continued)


	Task	Topic/Action Required	Primary Node P1 ¹	Standby Node P2
Step 9	Make sure the format of the /etc/hosts file is correct.	<p>Make sure the /etc/hosts file lists the hostname before the fully qualified domain name (FQDN).</p> <p>Bad /etc/hosts file:</p> <pre>127.0.0.1 localhost.localdomain localhost ::1 localhost6.localdomain6 localhost6 10.128.14.247 spin 172.16.17.127 cvldprimegate1.cscdev.com cvldprimegate1</pre> <p>Good /etc/hosts file:</p> <pre>127.0.0.1 localhost.localdomain localhost ::1 localhost6.localdomain6 localhost6 10.128.14.247 spin 172.16.17.127 cvldprimegate1 cvldprimegate1.cscdev.com</pre> <p>Also make sure the hostname is not mapped to the loopback address (localhost / 127.0.0.1).</p> <p> Note Make sure that the etc/host value and the system returned hostname value are the same. For example, in 172.16.17.127 cvldprimegate1 cvldprimegate1.cscdev.com the format for the Node 1 name should be specified as the second value of the hostname, that is cvldprimegate1.</p>	X	X
Step 10	Back up the /etc/host and root cron jobs files (the installation software will modify them).	—	X	X
Step 11	For cluster node makes sure the specified services are configured to start automatically each time the machine is rebooted.	Configure the Resources for Automatic Start After Reboot, page 3-13	X	X
Step 12	Stop the RHCS services in the order specified in Stopping the RHCS Services, page 3-14 .	Stopping the RHCS Services, page 3-14	X	X

Table 3-1 Steps for Setting Up and Installing Local Redundancy (continued)

	Task	Topic/Action Required	Primary Node P1 ¹	Standby Node P2
Step 13	Install the gateway and Oracle database using <code>install_prime_HA.pl</code> .	Installing the Prime Network Gateway Local Redundancy Software , page 3-15	x	—
Step 14	Configure the embedded database (using the <code>add_emdb_storage.pl -ha</code> script).			
Step 15	If desired, install any new device packages so that you have the latest device support.	Cisco Prime Network 5.2 Administrator Guide	x	x
Step 16	Verify the installation of the gateway and database.	Verifying the Local Redundancy Setup , page 3-21	x	x
Step 17	(Only for NAT) Update the database host.	Updating the Database Host in the Registry (Only for NAT) , page 3-24	x	—
Step 18	(Optional) Install Operations Reports.	Installing Prime Network Operations Report with Gateway High Availability	x	—
Step 19	(Optional) Install PN-IL.	Installing and Configuring PN-IL with Local Redundancy , page 3-30	x	—
Step 20	(Optional) Setup RHCS Web GUI if it is not configured during installation.	Configuring the RHCS Web Interface (Optional) , page 3-25	x	—

1. P1 is the primary cluster node and has the primary database.

Installation Requirements for Local Redundancy

These topics list the prerequisites for installing gateway geographical redundancy:

- [Hardware and Software Requirements for Local Redundancy](#), page 3-5
- [Port Usage for Local Redundancy](#), page 3-7

Hardware and Software Requirements for Local Redundancy

Table 3-2 shows the core system requirements for local redundancy. Local redundancy requires a Prime Network embedded database and does not support IPv6 gateways or databases. If your high availability deployment differs from these requirements, please contact your Cisco account representative for assistance with the planning and installation of high availability.


Table 3-2 Prerequisites for Local Redundancy ¹

Area	Requirements
Operating System	<p>Red Hat 6.7, Red Hat 6.8, Red Hat 6.9, Red Hat 6.10, Red Hat 7.4, or Red Hat 7.5, 64-bit Server Edition (English language). Red Hat can run in a virtual environment and supports VMware ESXi version 5.5, 6.0, or 6.7, and also on the Openstack kernel-based virtual machine (KVM) hypervisor version 2.6.</p> <p>Note Both nodes in the cluster must have identical RHCS versions and packages.</p> <p>Required Red Hat services and components:</p> <ul style="list-style-type: none"> • /usr/bin/expect—Tool to automate interactive applications • /usr/bin/ksh—Korn shell • /usr/bin/scp—Secure copy tool • /usr/sbin/sshd—SSH daemon • /usr/bin/ssh—SSH • /usr/bin/ssh-keygen—Tool to generate, manage, and convert authentication keys. <p>For more information on installing operating system and RPMs required on Red Hat, see Installing RHEL and Verifying the Version, page 3-9 and Installing RPMs Required on Red Hat for Prime Network, page 3-9.</p>
Oracle	<p>12.2.0.1.</p> <p>Note Oracle 12.2.0.1 is included in the Prime Network embedded database installation.</p>

Table 3-2 Prerequisites for Local Redundancy (continued)¹

Area	Requirements
Hardware	<p>RHEL 6.7, RHEL 6.8, RHEL 6.9, RHEL 6.10, RHEL 7.4, and RHEL 7.5 certified platform with fencing capabilities.</p> <p>Note RHEL supports the fence_vmware_soap fencing method on RHEL 6.5 or higher (with the High Availability and Resilient Storage Add Ons). For more information, see the Red Hat site. It is recommended for virtual machines, the RHCS must run with fence_vmware_soap fencing method.</p> <p>Note Hardware installation with no single point of failure is recommended. See Configuring Hardware and External Storage for Red Hat Cluster Site, page 3-8.</p> <p>While using fencing, ensure the following:</p> <ul style="list-style-type: none"> – Each node in the cluster uses a fencing method. – If you choose manual fencing option during the local redundancy installation to add a different Red Hat-supported fencing device, provision the device after installation using the RHCS GUI. When you add it, be sure to add it as the main fencing method, and move the manual fencing agent to the backup method. – 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 supply power to the node simultaneously. <p>Fencing options are listed in Fencing Options, page 2-3. For the recommended hardware for small, medium, and large networks, see the Cisco Prime Network 5.2 Installation Guide.</p>
Network	<ul style="list-style-type: none"> • Virtual IP Address <ul style="list-style-type: none"> – Reserve two floating IP addresses for ana and oracle_db services. These IP addresses are entered while executing the installation scripts. – Ensure that the IP addresses are on the same subnet and are not attached to any server. RHCS will manage them, that is, add and remove them from the server running the service. • Multicast Addresses <ul style="list-style-type: none"> – Cluster nodes must be able to communicate with each other using multicast. – Each network switch and 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. – Refer to the appropriate vendor documentation or other information about configuring network switches, and associated networking equipment, to enable multicast addresses and IGMP – Multicast address should meet RHCS requirements and should not be blocked by a firewall. If there is a firewall, disable it; see the Red Hat site for more information. – If you are using SELinux, it must be disabled (or in permissive mode). See the Red Hat site. • Network Timing Protocol (NTP) must be configured. For more details on procedures, see the Cisco Prime Network 5.2 Installation Guide.

Table 3-2 Prerequisites for Local Redundancy (continued)¹

Area	Requirements
Storage	<p>RHCS requires a shared storage accessible from all cluster nodes. When using external storage, ensure the following:</p> <ul style="list-style-type: none"> – All of the shared storage should have an ext3 file system installed. – Shared storage must be accessible and mountable from both nodes. The number of HDD, HDD types, HDD capacity, and RAID level, should be based on recommendations provided by the <i>Prime Network Capacity Planning Guide</i>. – Shared storage can be configured in several ways, it depends on your hardware. If there is only one link for the storage to the node, LABEL must be configured on each disk device. If the node is connected to the storage with more than 1 connection (recommended) multipath should be configured. – Each cluster service should use one partition. If the partitions are on the same disk, use a single partition for each service. If partitions are spread across disks, use a single disk for each service. Each disk must be labeled. <p> Note Labels used by any cluster service to name a distinct block device on any node in the cluster must be unique across all block devices. Also, a label used in a cluster service may not be reused by any block device with a different UUID, which is listed by command 'blkid', run the command on all nodes of the cluster, and cross check across all results, before configuring local HA cluster.</p> <p>If you are using Operations Reports, 1-4 additional partitions should be created for the Infobright database data, cache, backup, and DLP storage.</p>
File system	ext3
Disk space	5 GB under /tmp is required for installation

1. Virtual machine and bare metal requirements for hard disk, memory, and processor are same. Refer to the [Cisco Prime Network 5.2 Installation Guide](#) for memory and processor requirements.

Port Usage for Local Redundancy

In addition to the ports listed in the [Cisco Prime Network 5.2 Installation Guide](#), the following ports must be free.

You can check the status of the listed ports by executing the following command:

```
# netstat -tulnap | grep port-number
```

To free any ports, contact your system administrator.

Table 3-3 Additional Ports Required for Local Redundancy

Port No.	Used for:
9096	Operations Reports

Preparing to Install the Local Redundancy Solution

These topics describe the setup tasks you may need to perform before installing the local redundancy software:


- [Configuring Hardware and External Storage for Red Hat Cluster Site, page 3-8](#)
- [Installing RHEL and Verifying the Version, page 3-9](#)
- [Installing RPMs Required on Red Hat for Prime Network, page 3-9](#)
- [Configuring Disk Group and Volumes, page 3-11](#)
- [Verify That All Servers Are Ready for Installation, page 3-12](#)
- [Creating the Mount Points for Installation, page 3-13](#)
- [Stopping the RHCS Services, page 3-14](#)
- [Updating the Database Host in the Registry \(Only for NAT\), page 3-24](#)
- [Configuring the RHCS Web Interface \(Optional\), page 3-25](#)

Configuring Hardware and External Storage for Red Hat Cluster Site

Figure 3-1 shows the recommended hardware design to avoid a single point of failure, which includes:

- Disk mirroring at the storage location.
- Redundant RAID controllers.
- Redundant storage and gateway power supplies.
- Dual NICs on both gateways.
- Separate NIC connections to switches.
- NIC bonding in active/backup mode.

Figure 3-1 Local Redundancy Hardware Installation to Avoid Single Points of Failure



1	Dual NICs on both gateways	4	Redundant gateway power supplies
2	Disk mirroring	5	Redundant storage power supplies
3	Redundant RAID controllers		

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 shown in [Table 3-4](#) 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 3-4](#) can be created and configured to satisfy the Oracle performance requirements and protected with RAID5, RAID1, or RAID10. The Oracle volumes can be created on a single LUN.
- The number of HDD, HDD types, HDD capacity, and RAID level, should be based on recommendations provided by the *Prime Network Capacity Planning Guide*. Obtain the Capacity Planning Guide from your Cisco account representative.

Installing RHEL and Verifying the Version

Install the RHEL with the Red Hat Cluster Suite using the procedures in the Red Hat user documentation.

To verify that you have the required Linux version, use the following command:

```
cat /etc/redhat-release
```

RHEL installation version should be identical on all the servers.



Note

RHCS is included in the Red Hat Advanced Platform option. If Red Hat Clustering Service was not installed as part of RHEL, install the Red Hat Clustering Service using the procedures in the Red Hat user documentation.

Installing RPMs Required on Red Hat for Prime Network

These sections list the additional RPMs required for Red Hat and Oracle:

- [Required RPMs for Red Hat 6.5, page 3-9](#)
- [Required RPMs for Oracle Database 12c, page 3-10](#)

Required RPMs for Red Hat 6.5

The following RPMs must be downloaded from the Red Hat website and installed on the gateway and unit servers.

Required 32-bit packages

- compat-libstdc++-33-3.2.3-69.el6.i686
- glibc-2.12-1.132.el6.i686

- libgcc-4.4.7-4.el6.i686
- libstdc++-devel-4.4.7-4.el6.i686
- libaio-devel-0.3.107-10.el6.i686
- libXtst-1.2.1-2.el6.i686(Required for GUI installation)
- libgcj-4.4.7-4.1.el6_5.i686(Required for GUI installation)
- libXext.i686

Minimum Required 64-bit packages

- binutils-2.20.51.0.2-5.36.el6.x86_64
- libXtst-1.2.1-2.el6.x86_64 (Required for GUI installation)
- libgcj-4.4.7-4.1.el6_5.x86_64(Required for GUI installation)
- compat-libcap1-1.10-1.x86_64
- compat-libstdc++-33-3.2.3-69.el6.x86_64
- openssl098e-0.9.8e-17.el6_2.2.x86_64 (Required for installing Operations Reports)
- gcc-c++-4.4.7-4.el6.x86_64
- glibc-devel-2.12-1.132.el6_5.4.x86_64
- numactl-2.0.7-8.el6.x86_64
- ksh-20120801-10.el6.x86_64
- libgcc-4.4.7-4.el6.x86_64
- libstdc++-devel-4.4.7-4.el6.x86_64
- libaio-devel-0.3.107-10.el6.x86_64
- make-3.81-20.el6.x86_64
- sysstat-9.0.4-22.el6.x86_64
- expect-5.44.1.15-5.el6_4.x86_64
- openssh-server-5.3p1-94.el6.x86_64
- openssh-5.3p1-94.el6.x86_64
- telnet-0.17-47.el6_3.1.x86_64
- dos2unix-3.1-37.el6.x86_64

Required RPMs for Oracle Database 12c

The following packages, or later versions of them, are required for the Oracle 12c database on Red Hat.

- binutils-2.20.51.0.2-5.11.el6 (x86_64)
- glibc-2.12-1.7.el6 (x86_64)
- libgcc-4.4.4-13.el6 (x86_64)
- libstdc++-4.4.4-13.el6 (x86_64)
- libaio-0.3.107-10.el6 (x86_64)
- libXext-1.1 (x86_64)
- libXtst-1.0.99.2 (x86_64)

- libX11-1.3 (x86_64)
- libXau-1.0.5 (x86_64)
- libxcb-1.5 (x86_64)
- libXi-1.3 (x86_64)
- make-3.81-19.el6
- sysstat-9.0.4-11.el6 (x86_64)
- compat-libcap1-1.10-1 (x86_64)
- compat-libstdc++-33-3.2.3-69.el6 (x86_64)
- gcc-4.4.4-13.el6 (x86_64)
- gcc-c++-4.4.4-13.el6 (x86_64)
- glibc-devel-2.12-1.7.el6 (x86_64)
- ksh (any version of ksh)
- libstdc++-devel-4.4.4-13.el6 (x86_64)
- libaio-devel-0.3.107-10.el6 (x86_64)

**Note**

If any of the preceding packages are missing, the installation fails.

To verify all required RPMs are installed, execute the following command as root:

```
rpm -q binutils compat-libcap compat-libstdc++ expect gcc gcc-c++ glibc glibc-devel ksh
libgcc libstdc++ libstdc++-devel libaio libaio-devel make numactl numactl-devel sysstat
--qf' %{name} .%{arch}\n' | sort
```

Configuring Disk Group and Volumes

Table 3-4 and Table 3-5 show the disk partitions required for the dual-node cluster at the primary site.

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.
- Shared storage must be accessible from all cluster nodes. For recommendations on the number of HDD, HDD types, HDD capacity, and RAID level contact your Cisco representative.
- Placing the individual directories in separate partitions is recommended, though not required.

Table 3-4 Prime Network Local Redundancy Cluster Volume Sizes

Volume	Minimum Size (GB)	Comments
Prime Network	50	—
Oracle application + data files	10	—
Oracle redo logs	12.8	—
Oracle archives	20	See the <i>Prime Network Capacity Planning Guide</i> . Contact your Cisco account representative for information.

Table 3-4 Prime Network Local Redundancy Cluster Volume Sizes (continued)

Volume	Minimum Size (GB)	Comments
Oracle additional data files (if used)	—	Based on Prime Network alarm history needs. See the <i>Prime Network Capacity Planning Guide</i> .
Oracle backup	50	See the <i>Prime Network Capacity Planning Guide</i> .
If Operations Reports is installed: <ul style="list-style-type: none"> • Infobright data directory • Infobright cache directory • Infobright backup directory • Infobright DLP directory 	—	See the <i>Prime Network Capacity Planning Guide</i> and <i>Memory Assessment Tool</i> .

Table 3-5 Disk Groups

Partition	Space (in MB)
swap	Twice the size of the physical memory, up to 96 GB. For example, if your server has 16 GB RAM, the recommended swap space is 32 GB. If your server has 64 GB RAM, the recommended swap space is 96 GB.
/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 cluster conf

**Note**

Prime Network installation normally requires 1024 MB additional free space on the root partition. For HA, a temporary copy of Prime Network is installed under the root partition. Therefore, an additional 5120 MB free space is required, for a total of 6144 MB required free space. The HA files are installed under /usr/local/bin, /var, /etc., which requires a minimum of 1224 MB. You can add this amount to the root partition instead of creating a separate partition for each.

Verify That All Servers Are Ready for Installation

Verify the following on all servers: disk access, Linux versions, and NTP sync on all servers:

- Access to all external disks is available.
- The same version of Linux is deployed on all servers. To check the version:

```
cat /etc/redhat-release
```

- Verify that the time is synchronized on both servers using NTP. For information on configuring NTP, see the [Cisco Prime Network 5.2 Installation Guide](#).

Creating the Mount Points for Installation

Use this procedure to create mount points before setting up high availability.



Note

All servers in the local redundancy setup should have same mount points.

- Step 1** Log in as root user, and create the following directories:
- Prime Network home directory and Oracle directories.
- ```
mkdir -p /pn41
mkdir -p /opt/ora
mkdir -p /redo
mkdir -p /data
```
- Operations Reports directories (applicable for Operations Reports).
- ```
mkdir -p /ldata
mkdir -p /lcache
mkdir -p /lbackup
mkdir -p /ldlp
```
- Step 2** Mount the external shared storage on the relevant directories of the node from where you will run the installation. Mount it manually and do not add it to the fstab file. Comment out any corresponding entry to the shared storage in /etc/fstab for both cluster nodes.
- Step 3** If the embedded database mount points contained in networkdata/archive logs and control files are set outside the local disks, for example, on a SAN, make corresponding entries in /etc/fstab so the mount points are available during a reboot.
- Step 4** Mount all of the Oracle, Prime Network, Operations Reports mount points on the server where you will run the installation.

In this example, PRIMENETWORK and ORACLE are the sample label names:

```
mount -L PRIMENETWORK/pn41
mount -L ORACLE/opt/ora
mount -L PRIMENETWORK/redo
mount -L PRIMENETWORK/data
```

```
mount /dev/sda1 /ldata
mount /dev/sda2 /lcache
mount /dev/sda3 /lbackup
mount /dev/sda4 /ldlp
```

Configure the Resources for Automatic Start After Reboot

For every cluster node, make sure the following resources are configured to start automatically each time the server is rebooted.

- modclusterd

- ricci
- rgmanager
- cman

For automatically starting these services, run the following command:

```
chkconfig modclusterd on
chkconfig ricci on
chkconfig rgmanager on
chkconfig cman on
```

Check that status of these services using the following command:

```
chkconfig --list ricci
ricci 0:off      1:off      2:off      3:off      4:off      5:off      6:off
```

The above output indicate the ricci service is disabled

```
chkconfig --list ricci
sshd 0:off      1:off      2:on       3:on       4:on       5:on       6:off
```

The above output indicate the ricci service is enabled

Stopping the RHCS Services

Make sure that the Red Hat Cluster Suite rgmanager and cman services are turned off before installing Prime Network high availability on the gateway.

To turn off the RHCS services:

Step 1 On P1, stop the rgmanager service using the following command:

```
service rgmanager stop
```

Step 2 On P2, stop the rgmanager service using the following command:

```
service rgmanager stop
```

Step 3 On P1, stop the cman service using the following command:

```
service cman stop
```

Step 4 On P2, stop the cman service using the following command:

```
service cman stop
```

Step 5 Enter the following command on all cluster nodes to verify the service status:

```
service rgmanager status
service cman status
```

Step 6 The services are stopped.

For rgmanager stopped services the output is displayed as `clurgmgrd is stopped` and for cman as `ccsd is stopped`.

Installing the Prime Network Gateway Local Redundancy Software

The local redundancy solution for dual-node cluster is installed using **install_prime_HA.pl** script that is available in RH_ha.zip file in the installation DVD as described in [Installation DVDs, page 1-1](#).

You can run the installation in interactive or in non-interactive mode. Interactive mode installation prompts you to enter the gateway HA data values one at a time. The Prime Network installer then updates the auto_install_RH.ini file template, which populates the **install_Prime_HA.pl** script.

Alternatively, you can enter all the installation values in the auto_install_RH.ini template, located in the RH_ha directory, then run the installation in non-interactive mode. The installation mode is determined by the presence or absence of the **-autoconf** flag.



Note

It is recommended you run the installation in interactive mode first to populate the auto_install_RH.ini template with the user input. This gives you the ability to verify the input and run the installation again in non-interactive mode, if needed.

This procedure installs gateway high availability for local redundancy with RHEL 6.5, 6.7, or 6.8.

- Step 1** Change to the root user, then unzip the RH_ha.zip located on the installation DVD. Unzipping RH_ha.zip creates the /tmp/RH_ha directory.



Note If you are running the Korn shell (/bin/ksh) and the prompt is the hash tag (#), the installation will fail. Run the installation script using bash.

- Step 2** From the /tmp/RH_ha directory run the **install_Prime_HA.pl** in interactive or non-interactive mode. For information on the **install_Prime_HA.pl** script, see [Installation DVDs, page 1-1](#).
- Step 3** For local redundancy *alone*, enter local HA= yes, DR= no, when prompted. See [Table 3-6](#) for the prompts that appears while installing local redundancy configuration.
- Step 4** Execute the **install_Prime_HA.pl** script in interactive or non-interactive method.

- **Interactive Installation:**

For interactive installation, execute the following commands:

```
cd /tmp/RH_ha
perl install_Prime_HA.pl
```

See [Table 3-6](#) for descriptions of other parameters you will be asked to enter at various stages of the interactive installation.

- **Non-Interactive Installation (Automatic):**

- a. Edit the auto_install_RH.ini file template found under the RH_ha directory with all of the installation details.
- b. Run the following command:

```
cd /tmp/RH_ha
perl install_Prime_HA.pl -autoconf <full-path-of-auto_install_RH.ini-file>
```



Note To prevent any security violation, it is highly recommended to remove the password in `auto_install_RH.ini` file after the successful installation.

After the `install_Prime_HA.pl` script is completed:

- Prime Network and embedded database will be installed on the setup node. The cluster standby node will have only the users and home directory.
- RHCS will be up and running the Prime Network (`ana`) and Oracle (`oracle_db`) services.

[Table 3-6](#) describes the prompts that you need to enter during the local redundancy installation.



Note If you experience problems, see [Troubleshooting the Local Redundancy Installation, page 3-20](#).

Table 3-6 Installation Prompts for Local Redundancy Alone

Prompt for	Enter...	Notes
Configure local HA?	yes	—
Configure DR?	no	Enter no ; this procedure is for local redundancy <i>alone</i> . To install local + geographical redundancy, see Installing the Prime Network Gateway Geographical Redundancy Software, page 4-6 .
Configure NTP on 2 gateways?	yes	yes or no depending on whether NTP should be configured on two gateways. If not configured, first configure NTP and then continue with the installation. For more details on procedures, see configuring NTP in the Cisco Prime Network 5.2 Installation Guide .
OS user of the database	oracledb	Oracle installation owner (default is oracle).
Prime Network OS user	<i>pnuser</i>	User-defined Prime Network OS user (<i>pnuser</i>). Username must start with a letter and contain only the following characters: [A-Z a-z 0-9].
Oracle user home directory	Home directory of user oracle	Location of the mount point given for the <i>oracle-home/oracle-user</i> . Default is <code>/opt/ora/oracle</code> .
Home directory of the Prime Network user	Example: <code>/export/home/ana/pn41</code>	Directory should be located under <i>Prime Network file system mount point</i> but <i>not</i> the mount point itself.
Prime Network user password	<i>password</i>	User-defined password for the <i>pnuser</i> .
Location of the Prime Network installation file	Example: <code>/dvd/Server</code>	Mount point of the Prime Network installation. Should be the same for all relevant nodes. Example: For install.pl the path will be <code>/dvd/Server</code> .

Table 3-6 Installation Prompts for Local Redundancy Alone (continued)

Prompt for	Enter...	Notes
Oracle mount point	Example: /opt/ora	Location of Oracle mount points, separated by ",". First is the mount point for the Oracle home directory, for example, /opt/ora,/opt/dbf. Note For interactive installations: Installer asks you for a mount, then asks if you want to add another one. For non-interactive installations, enter all Oracle mount data in the input file.
Configure another oracle file system mount	no	yes or no value indicating whether you want to use the default Oracle mount point or not
Prime Network mount point	Example: /export/home	Location of Prime Network mount point.
Directory for the Oracle zip files	Example: /opt/ora/oracle_zip	Directory containing embedded Oracle zip files. Can be a temporary location where the files were copied from the installation DVDs; or directly specify the location on DVD.
Node one name	node 1 hostname	Hostname for node running the installation. For local redundancy dual-node clusters, node must be one of the cluster nodes. This is the value returned by the system call hostname.
Node two name	node 2 hostname	hostname for the second cluster node for local redundancy dual-node clusters. This is the value returned by the system call hostname.
DB profile	The number corresponding to the DB profile required.	Select from (1-7). Estimated DB profile.
Password for 5 built-in users	password	Password for Prime Network root, bosenable, bosconfig, bosusermgr, and web monitoring users (users for various system components). Passwords must contain: <ul style="list-style-type: none"> • Contain at least eight alphanumeric characters. • Contain upper and lower case letters. • Contain one number and one special character. • Cannot contain: @ / ! \$ ~ * () - + = [{
Running database backups.	yes/no	Whether to enable embedded database automated backups.
SMTP server	Example: outbound.cisco.com	Local e-mail server.

Table 3-6 Installation Prompts for Local Redundancy Alone (continued)

Prompt for	Enter...	Notes
User email	email address	<p>E-mail address to which embedded database will send error messages.</p> <p>When a local HA Oracle database is switched to run on a different gateway either manually or automatically, the oracle started in machine_name notification will be emailed to the recipient with email address configured in oracle.sh. If you want a different recipient to receive the email notification, you need to manually update the oracle.sh file.</p> <p>For example,</p> <pre>[root@pslucbpngd1 ~]# less /usr/local/bin/oracle.sh #!/bin/bash # Global variables ORACLE_USER=oracle HOMEDIR=/oracle/oracle ORACLE_MOUNT1=/oracle ORACLE_MOUNT2=/oradata ORACLE_MOUNT3=/redo01 ORACLE_MOUNT4=/archduplex OVERRIDE_FILE=/var/tmp/override REC_LIST= jpratap@cisco.com</pre>
DB archive	Example:/opt/ora/oracle/arch	Location of the database archive files. Should be located under one of the Oracle mounts but not directly on the mount, and should be compliant with the storage requirements.
DB redo	Example: /opt/ora/oracle/redo	Location of database redologs. Should be located under one of the Oracle mounts but not directly on the mount, and should be compliant with the storage requirements.
DB backup dest	Example:/opt/ora/oracle/back up	Location of database backup files. Should be located under one of the Oracle mounts but not directly on the mount, and should be compliant with the storage requirements.
DB datafiles	Example:/opt/ora/oracle/data	Location of database data files. Should be located under one of the Oracle mounts but not directly on the mount, and should be compliant with the storage requirements.
Oracle service IP address	IP address	Virtual IP of local cluster Oracle service group.
Prime Network service IP address	IP address	Virtual IP of local cluster Prime Network service group.
Multicast address	IP address	An available multicast address accessible and configured for both cluster nodes.
Prime Network cluster name	<i>username</i>	User-defined cluster name. The cluster name cannot be more than 15 non-NUL (ASCII 0) characters. For local HA, the cluster name must be unique within the LAN.
Node one fence agent	The number corresponding to the fencing agent required	Type of fencing device configured for node running the installation. (For information about supported fencing devices and information you may need to provide to the installation script, see supported fencing methods in Fencing Options, page 2-3.)

Table 3-6 Installation Prompts for Local Redundancy Alone (continued)

Prompt for	Enter...	Notes
Node one fence hostname	hostname	Hostname of fencing device configured for the node running the installation (for some fencing devices, can be an IP address).
Node one fence login	login name	Login name for fencing device configured for node running the installation.
Node one fence passwd	password	Password for fencing device configured for node running the installation.
Node two fence agent	The number corresponding to the fencing agent required	Type of fencing device configured for second cluster node. (For information about supported fencing devices and information you may need to provide to the installation script, see supported fencing methods in Fencing Options, page 2-3.)
Node two fence hostname	hostname	Hostname of fencing device configured for second cluster node (for some fencing devices, can be an IP address).
Node two fence login	login name	Login name for fencing device configured for second cluster node.
Node two fence passwd	password	Password for fencing device configured for second cluster node.
Prime Network cluster web interface password	port number and password	ort and the password for cluster web interface. <i>LUCI_PORT</i> must be available and should not be in Prime Network debug range: 60000 <= <i>x</i> < 61000 or in Prime Network AVM port range: 2000 <= <i>x</i> < 3000 or 8000 <= <i>x</i> < 9000 Password must contain at least 6 characters.
Prime Network cluster web interface port		

Step 5 Configure the embedded database by running the **add_emdb_storage.pl** utility. In the following, *NETWORKHOME* is the Prime Network installation directory (*/export/home/pnuser* by default).

- a. Log in as *pnuser*.

```
su - pnuser
```

- b. Change directories to *NETWORKHOME/Main/scripts/embedded_db* and enter the following command:

```
perl add_emdb_storage.pl -ha
```

Enter the number corresponding to the estimated database profile that meets your requirement. For more information, contact your Cisco representative and obtain the *Prime Network Capacity Planning Guide*.

- c. Insert the event and workflow archiving size in days. If you are not sure what to choose, take the default.

When you are done, validate the installation by following the steps in [Verifying the Local Redundancy Setup, page 3-21](#).

Troubleshooting the Local Redundancy Installation

1. Problem: Installation fails.

Solution: Please review the following:

- Make sure all the necessary ports for installation are free, otherwise installation prerequisite verification returns an error that a needed port is blocked.
- For a virtual machine, if the installation prerequisite verification returns an error that swap space is insufficient, you can override the message and continue the installation by adding the following entry into the `auto_install_RH.ini` file.

```
OVERRIDE_SWAP=true
```



Note Changing the Override Swap value to True is not recommended because a Prime Network service might not function correctly without the required swap space.

- If the failure occurs because a parameter needs to change, save the `auto_install_RH.ini` file to a temporary directory, then remove the old `RH_ha` directory and files. After you remove the old directory and files, redeploy the **RH_ha.zip** file. You must do this because installation changes the template files. However, after correcting the incorrect parameters, you can use the old `auto_install_RH.ini` file so you do not have to enter the correct input parameters again.
- If a local service (network/oracle_db services) in a local redundancy configuration fails, RHCS will try to stop, unmount, mount, then start the service locally. If this does not succeed, RHCS will automatically try to relocate the service to the standby node.
- If the local redundancy cluster nodes lose connection to each other, they try to fence each other. The node that succeeds starts the cluster services.
- If a local redundancy service enters a stopped state and does not start automatically on either node, you can start the service using the RHCS web or CLI interface. Before you do this, review the cluster log located in the `/var/log/messages`.
- When you run the **install_Prime_HA.pl** script log files are created. These are located in `tmp/RH_ha`.
- If the Prime Network file replication (not the Oracle database) in a geographical redundant configuration fails, verify the root cron jobs on both the primary and remote sites. The cron list and scripts run by the crons are located in the `/var/adm/cisco/prime-network/scripts/ha/rsync` directory.



Note

If you need to reinstall an embedded database in a directory that previously contained an embedded database, you must manually remove the database. If you do not do this, the installation will fail.

2. Problem: PN service relocation fails.

Description: This occurs when the unmount operation fails.

Solution: Cisco Advanced Malware Protection (AMP) service has to be disabled to resolve this issue. Follow the below mentioned steps:

1. Check the status of Cisco AMP service. Execute the following command:

On RHEL 7.x:

```
systemctl status cisco-amp
```

On RHEL 6.x:

```
initctl status cisco-amp
```

- If Cisco AMP is enabled, execute the following commands to disable it:

On RHEL 7.x:

```
cd/opt/cisco/amp/bin
systemctl disable cisco-amp
systemctl stop cisco-amp
```

On RHEL 6.x:

```
cd/opt/cisco/amp/bin
initctl stop cisco-ampmon
```

- Perform service relocation.
- Re-enable Cisco AMP service after relocating the services. Execute the following commands:

On RHEL 7.x:

```
systemctl start cisco-amp
```

On RHEL 6.x:

```
initctl start cisco-ampmon
```

- Repeat Step 1 to verify the status of Cisco AMP service.

Verifying the Local Redundancy Setup

To verify the installation, perform the verification steps in [Table 3-7](#). After you have verified the setup, proceed to [Post-Installation Tasks for Local Redundancy](#), page 3-24.

Table 3-7 Local Redundancy Verification Tests

Description	Procedure	Expected Results
<i>Local Cluster Hardware Failure</i>		
<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"> Power off the active node that runs both services (Prime Network and DB). 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>
<i>Manual Cluster Administration</i>		
<p>Name: Manual Service Stop</p> <p>Purpose: Verify that the service can be manually stopped.</p>	<ol style="list-style-type: none"> Enter: <pre>clusvcadm -d service_name</pre> Verify 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 manually started.</p>	<ol style="list-style-type: none"> Run <pre>clusvcadm -e service_name</pre> Verify that it is running and no errors exist in cluster log (/var/log/messages on both cluster nodes). 	<p>The service is running.</p>

Table 3-7 Local Redundancy Verification Tests (continued)

Description	Procedure	Expected Results
Name: Manual Service Relocation Purpose: Verify that the service can be manually relocated.	<ol style="list-style-type: none"> Enter: <code>clusvcadm -r service_name</code> Verify that the service is not running on the current node and is running on the standby node. 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. Test both the Prime Network and Oracle services. 	The service is stopped on the active node and started on the redundant node.
<i>Ordered Cluster Node Startups</i>		
Name: Node Startup in Existing Cluster Purpose: Verify that a cluster node starts up and rejoins a cluster after it is restarted.	<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.
Name: Simultaneous Node Startup Purpose: Verify that the cluster is set up correctly when both nodes start simultaneously.	<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 are running on the cluster. Check the log for errors. 	Both cluster nodes join the cluster; both services are running.
Name: Single Node Startup Purpose: Test the cluster functionality when only one is node running.	<ol style="list-style-type: none"> Power down both nodes, then start one of them. The running node will fence the other node and run the services. The fenced node joins the cluster to create the dual node cluster. Check log for errors. 	Both cluster nodes join the cluster; both services are running.
<i>Local Cluster Service Failure</i>		
Name: Service Failure Purpose: Test the service startup after a failure occurs.	<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 it was running. Check the log for errors. Test both the Prime Network and Oracle services. 	The service is restarted on the same node.
<i>Local Cluster HW Failure</i>		

Table 3-7 Local Redundancy Verification Tests (continued)

Description	Procedure	Expected Results
<p>Name: Stop Node with Fencing Off</p> <p>Purpose: Verify the node requires manual fencing after the other node, including its fencing agent, is removed.</p>	<ol style="list-style-type: none"> 1. Disconnect the fencing agent to one of the nodes, then power it off unexpectedly. 2. Observe the other node behavior. 3. Check the log for errors and the request for manual fencing. 	<p>The fence_ack_manual required notification appears in the logs. A message is printed to /var/log/messages advising you to run the fence_ack_manual command on the gateway server.</p> <p>The cluster is running with one node and all services running on it.</p>
<p>Name: Single Node Cluster</p> <p>Purpose: Checks that the cluster can function when the other node does not exist at all, or have no power at all</p>	<ol style="list-style-type: none"> 1. Power down both nodes, 2. Disconnect the fencing agent to one of the nodes. 3. Start the other node. It will attempt to fence the other node, but fail with the regular fencing agent. Manual fencing is required. 4. Acknowledge the manual fencing. 	<p>The cluster does not start the services (and does not show in the clustat command) before acknowledging that manual fencing is performed.</p>

Verifying Cluster Services

<p>Name: Verifying Cluster Services</p> <p>Purpose: Verify that the cman and rgmanager services are running on both cluster nodes.</p>	<p>As the root user, enter the following command to verify the cluster and services are running:</p> <pre>clustat</pre> <p>Example:</p> <pre>[root@hostname RH_ha] clustat Cluster Status for network_cluster @ Mon Apr 16 10:01:02 2013 Member Status: Quorate Member Name ID Status ----- ----- hostname.cisco.com 1 Online, Local, rgmanager hostname2.cisco.com 2 Online, rgmanager Service Name Owner (Last) State service:ana hostname.cisco.com started service:oracle_db hostname.cisco.com started</pre>	<p>Verify that the cman and rgmanager services are running on both cluster nodes.</p>
--	---	---

Post-Installation Tasks for Local Redundancy

After you have validated the installation, perform these post-installation tasks:

- [Updating the Database Host in the Registry \(Only for NAT\), page 3-24](#)
- [Configuring the RHCS Web Interface \(Optional\), page 3-25](#)

Updating the Database Host in the Registry (Only for NAT)

If you are using network address translation (NAT) with the Cisco Prime Network Vision client, update the database host in the Prime Network registry to contain the hostname instead of the IP address.

Complete the following mandatory steps after the Cisco Prime Network 5.2 gateway installation or upgrade is complete and the system is up and running.



Note

If you already use a hostname instead of an IP address, you do not have to repeat this procedure.

In the following procedure, *NETWORKHOME* is the Prime Network installation directory (*/export/home/pnuser* by default).

Step 1 Before changing the hostname, verify that the Windows client workstations have the correct Domain Name System (DNS) mapping.

Step 2 From *NETWORKHOME/Main*, enter the following commands:

```
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 site/persistency/nodes/main/Host
database-server-hostname
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 site/persistency/nodes/ep/Host
database-server-hostname
```

During switchover, you should unset the entries in the *site.xml* file and then reset using the following commands:

```
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/main/Host
database-server-hostname
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/ep/Host
database-server-hostname
```

You can also change the FQDN in all nodes of *persistency.xml*.

Example:

```
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/infobright/Host
database-server-hostname
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/ep_rep/Host
database-server-hostname
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/main_rep/Host
database-server-hostname
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/admin/Host
database-server-hostname
./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 persistency/nodes/xmp/Host
database-server-hostname
```

Step 3 Enter the following command to restart the Prime Network system:

```
networkctl restart
```


Configuring the RHCS Web Interface (Optional)

The RHCS web interface is configured during the install process. Use the information provided in this section only if you decide to change the configuration of the web interface at a later stage or if the web interface was not configured during the installation process.

The RHCS “luci” web interface allows you to configure and manage storage and cluster behavior on remote systems. You will use it to manage the Prime Network gateway HA. Before you begin this procedure, you should have the Red Hat *Conga User Manual*. It can be obtained at:

http://sources.redhat.com/cluster/conga/doc/user_manual.html

If your fencing device is supported by RHCS but not listed in [Fencing Options, page 2-3](#), that is, you chose the Manual fencing option during the installation, manually configure the device using the Red Hat fencing configuration documentation. This can be obtained at:

http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/pdf/Configuration_Example_-_Fence_Devices/Red_Hat_Enterprise_Linux-5-Configuration_Example_-_Fence_Devices-en-US.pdf



Note

The following procedure provides the general steps to configure the luci interface. See the Red Hat *Conga User Manual* for details on performing steps in this procedure.



Note

The RHCS web interface must be configured for both servers in the local redundant dual-node cluster.

Step 1 As root user, run the following command and enter the needed details:

```
luci_admin init
```

Step 2 Edit `/etc/sysconfig/luci` to change the default port to an available port. (The default 8084 port is used by Prime Network.) For example:

```
defaults for luci,
web UI fronted for remote cluster and storage management
```

```
LUCI_HTTPS_PORT=8085
```



Note

The ports must be available and should not be in Prime Network debug range (60000 <= X < 61000) or in Prime Network avm port range (2000 <= X < 3000 or 8000 <= X < 9000)

Step 3 As the root user, enter:

```
service luci restart
```

Step 4 Enter the web interface using the following link:

```
https://node host name:port
```

From the RHCS web interface you can stop, start, and relocate the Prime Network and oracle_db services managed by the cluster.

- Step 5** In the luci web interface, add the cluster that was configured by the Prime Network installation. See the Red Hat *Conga User Manual* for details on performing the following:
- Add a system.
 - Add an existing cluster.
 - Add a user.
- Step 6** If your fencing device is supported by RHCS but not listed in [Fencing Options, page 2-3](#), use the Red Hat fencing configuration guide to configure the device.



Note If you provision a new fencing device, set it as the primary fencing method. The manual fencing agent should be kept as the backup fencing method.

Maintaining Local Redundancy

After the local redundancy cluster is deployed, failovers are automatic. In case of a single service failure, the cluster will attempt to restart the service. If the retries fail, the service will be relocated to the second node and started on that node. This does not impact the other service in the cluster.



Note For complete redundancy, a configuration with no single point of failure is recommended. See the RHCS documentation for recommended configurations.

- [Monitoring Log Messages, page 3-26](#)
- [Monitoring Cluster Status Using the CLI, page 3-27](#)
- [Monitoring Cluster Status Using the GUI, page 3-27](#)
- [Managing the Local Redundancy Cluster, page 3-27](#)
- [Manually Fencing, page 3-29](#)

Monitoring Log Messages

The RHCS log messages provide information about cluster-related issues, such as service failure.

Every 30 seconds, RHCS issues status commands to check the Prime Network, Oracle, and Oracle listener processes. These messages are logged to `/var/log/messages` and can be viewed by the root user (or from the RHCS web GUI). The following are some example messages.

```
Mar 23 13:45:47 hostname clurgmgrd: [27961]: <info> Executing /usr/local/bin/ana.sh status
Mar 23 13:46:07 hostname clurgmgrd: [27961]: <info> Executing /usr/local/bin/oracle.sh
status
Mar 23 13:46:07 hostname clurgmgrd: [27961]: <info> Executing /usr/local/bin/lsnr.sh
status
```

Monitoring Cluster Status Using the CLI

You can use the **clustat** command checks a cluster's members and overall status.

As the root user, enter the following command to verify the cluster and services are running:

```
clustat
```

In the following example, the cluster name is `ana_cluster` and `hostname.cisco.com` is the node from which the command was run.

```
root@hostname.cisco.com] clustat
Cluster Status for ana_cluster @ Thu Mar  3 10:24:50 2014
Member Status: Quorate

Member Name                ID      Status
-----
hostname.cisco.com         1      Online, Local, rgmanager
hostname2.cisco.com        2      Online, rgmanager

Service Name                Owner (Last)      State
-----
service:ana                 hostname.cisco.com  started
service:oracle_db           hostname2.cisco.com  started
```

Monitoring Cluster Status Using the GUI

The RHCS web interface is automatically configured by the Prime Network installation script. If the interface was not configured during the installation process, use the procedure in [Configuring the RHCS Web Interface \(Optional\)](#) section to configure RHCS Web GUI. For details on how to use the web GUI, see the appropriate RHCS documentation.

Web GUI is used to:

- Check the cluster status, including the status of each service and the node each service is running on.
- Initiate a switchover of a service to the other node (relocate the service from the Services area of the GUI).

You can connect to the RHCS web interface by entering the following in the address field of your browser: **https://cluster-node-hostname:port/luci**.

Managing the Local Redundancy Cluster

You can use `clusvcadm` command to check the version of the RHCS used on the cluster, stop, restart the cluster services and so on. To manage the cluster from the CLI, enter:

```
[root@hostname RH_ha] clusvcadm
```

[Table 3-8](#) shows the RHCS `clusvcadm` command options to manage the cluster.

Table 3-8 RHCS CLI Commands

<code>clusvcadm +</code>	Description
<code>-v</code>	Display version and exit
<code>-d group</code>	Disable <i>group</i>

Table 3-8 RHCS CLI Commands (continued)

clusvcadm +	Description
<code>-e group</code>	Enable <i>group</i>
<code>-e group -F</code>	Enable <i>group</i> according to failover domain rules
<code>-e group -m member</code>	Enable <i>group</i> on <i>member</i>
<code>-r group -m member</code>	Relocate <i>group</i> to <i>member</i>
<code>-M group -m member</code>	Migrate <i>group</i> to <i>member</i> (e.g. for live migration of VMs)
<code>-R group</code>	Restart a <i>group</i> in place
<code>-s group</code>	Stop <i>group</i>
<code>-Z</code>	Freeze <i>group</i> in place
<code>-U</code>	Unfreeze/thaw <i>group</i>

To restart Prime Network, Oracle, or (if installed) Operations Reports application processes, use the following procedure.

- Step 1** Place the Prime Network and database RHCS services in maintenance mode (also called freezing) using the following command, where *service* is **ana**, **oracle_db**, or **ifb**.

```
clusvcadm -Z service
```

- Step 2** Confirm that the services are in maintenance mode. Run **clustat** and verify that the output shows the service followed by a [Z], which indicates the service is in maintenance mode (frozen). When the services are frozen, the cluster does not monitor them.

```
root@hostname.cisco.com] clustat
Cluster Status for ana_cluster @ Thu Mar  3 12:31:55 2013
Member Status: Quorate

Member Name                ID          Status
-----
hostname.cisco.com         1          Online, rgmanager
hostname.cisco.com         2          Online, Local, rgmanager

Service Name                Owner (Last)      State
-----
service:ana                 hostname.cisco.com started [Z]
service:oracle_db           hostname.cisco.com started [Z]
```

**Note**

If you attempt to restart either the Prime Network, Oracle, or Infobright applications without freezing the RHCS process, the cluster may detect that the services are down and attempt to restart them.

- Step 3** After confirming that the **ana**, **oracle_db**, and **ifb** cluster configured services are frozen, use the normal application commands to stop Prime Network and Oracle.

```
clusvcadm -s group
```

- Step 4** After restarting the Prime Network, Oracle, and Infobright applications, move the RHCS services out of freeze mode and reinitiate the cluster's monitoring of the ana and oracle services:

```
clusvcadm -U group
```

Manually Fencing

During the installation of the RHCS solution, you are prompted to select one of three fencing options. You can reconfigure the fencing choice at any time using the RHCS web interface or other RHCS tools. If you choose manual fencing, you must disconnect the node and storage when a problem occurs (either by disconnecting the node and storage by hand or by using another fencing agent).

**Note**

We recommend that manual fencing only be used on a temporary basis. If you use manual fencing, it is your responsibility to make sure that when an error occurs, the node and the storage are disconnected during the cluster workflow. We recommend that manual fencing only be used on a temporary basis and as a backup for your chosen fencing agent.

If you are using manual fencing and an error occurs that requires fencing intervention, a message is printed to `/var/log/messages` advising you to run the **fence_ack_manual** command on the gateway server.

**Note**

(Only for Red Hat 6.x)

- Before disconnecting the faulty node, remove the **cman** and **rgmanager** services from the automatic startup sequence. This is to avoid the failure when the restored node joins the cluster. You can remove these services by using the commands,

```
chkconfig -del cmanan
chkconfig -del rgmanager
```

- Start these services after the servers are restored using the following command:

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

Use the procedure below to disconnect the faulty node.

Step 1 Log into the gateway server as root and enter the command using the following syntax.

```
fence_ack_manual -n nodename
```

where **n** *nodename* indicates the node that has been disconnected from storage.

**Note**

For Red Hat 6.x, use only **fence_ack_manual** *nodename*.

Step 2 Continue with the confirmation message to disconnect the faulty node from the storage.

Uninstalling Local Redundancy

To uninstall local redundancy setup, follow the procedure provided below. The procedure also removes the operations reports if installed.

Step 1 Determine the active and standby cluster nodes using the following command:

clustat

Step 2 Stop cluster services on the standby nodes using the following command:

service regmanager stop

service cman stop

Step 3 Uninstall the Prime Network on the standby nodes by choosing **Yes** to all the prompts. The process uninstalls the Prime Network even if the disks are not mounted on the secondary nodes.

/var/adm/cisco/prime-network/reg/current/uninstall.pl

Step 4 Verify if the configuration file (cluster.conf) located at **/etc/cluster/cluster.conf** is either deleted or rolled back to the state before the Prime Network is installed.

Step 5 Stop the cluster services on the active node using the following commands:

service regmanager stop

service cman stop



Note

Verify if ana/oracle disks are not mounted by using the command **df -h**

Step 6 Manually remount the filesystems used by the cluster on the correct mounting points. This is because shutting down the cluster results in dismounting of all filesystems related to the cluster services.

Step 7 Uninstall Prime Network on the active node using the following command.

/var/adm/cisco/prime-network/reg/current/uninstall.pl

Step 8 Verify if the configuration file (cluster.conf) located at **/etc/cluster/cluster.conf** is either deleted or rolled back to the state before the Prime Network is installed.

Installing and Configuring PN-IL with Local Redundancy

This section explains how to install and configure the Prime Network Integration Layer (PN-IL) 1.2 with a Prime Network gateway local redundancy deployment. It also explains how to integrate the deployment with Cisco Prime Central. For information on the Prime Central releases with which you can install PN-IL 1.2, see the [Cisco Prime Network 5.2 Release Notes](#).

These topics provide the information you will need to install and configure PN-IL local redundancy:

- [Installation DVD](#), page 3-31
- [Steps for Installing PN-IL with Local Redundancy](#), page 3-31
- [Installing PN-IL on a Prime Network Server \(Local Redundancy\)](#), page 3-32
- [Configuring PN-IL on a Prime Network Gateway \(Local Redundancy\)](#), page 3-33
- [Disabling the PN-IL Health Monitor](#), page 3-36

If you want to migrate an *existing* standalone installation of PN-IL (with local redundancy) to suite mode, you can use the procedure in [Configuring PN-IL with Prime Central \(Suite Mode with Local Redundancy\)](#), page 3-34.

Installation DVD

The PN-IL high availability files are provided on the Prime Network installation DVD named **Disk 1: New Install DVD**. **Disk 2** contains the tar file **sil-esb-2.2.0.tar.gz**, which contains the PN-IL installation files and scripts, including:

- `installAndConfigureESB.sh`—PN-IL installation script
- `itgctl`—PN-IL configuration script
- `il-watch-dog.sh`—PN-IL health monitor control script
- `DMSwitchToSuite.sh`—Script to migrate to suite

Steps for Installing PN-IL with Local Redundancy

[Table 3-9](#) provides the basic steps you must follow to set up local redundancy for PN-IL.



Note

Install PN-IL only on the primary server.

If you want to migrate an *existing* standalone installations of PN-IL (with local redundancy) to suite mode, see the procedure in [Configuring PN-IL with Prime Central \(Suite Mode with Local Redundancy\)](#), page 3-34.

Table 3-9 Steps for Installing PN-IL Local Redundancy

	Task	Topic/Action Required	Server (P1) (has Primary database)	Server (P2)
Step 1	Collect server details, so that you have all information handy prior to installation.	<ul style="list-style-type: none"> • Virtual IP address of P1 • Prime Network application root username and password for P1 • URL for authenticating Prime Network calls on P1 (normally https://localhost:6081/ana/services/userman) • (Suite mode) For the Prime Central server where Oracle is installed: Hostname, database service name, database username and password, and database port. 	x	—
Step 2	Verify the server meets the prerequisites.	Installation Requirements for Local Redundancy , page 3-4	x	—
Step 3	Freeze RHCS and install PN-IL.	Installing PN-IL on a Prime Network Server (Local Redundancy) , page 3-32	x	—

Table 3-9 Steps for Installing PN-IL Local Redundancy (continued)

	Task	Topic/Action Required	Server (P1) (has Primary database)	Server (P2)
Step 4	Configure PN-IL (in standalone or suite mode) and unfreeze RHCS.	Configuring PN-IL on a Prime Network Gateway (Local Redundancy), page 3-33	x	—
Step 5	Disable the PN-IL Health Monitor	Disabling the PN-IL Health Monitor, page 3-36	x	—

Installing PN-IL on a Prime Network Server (Local Redundancy)

Before You Begin:

Make sure Prime Network is installed and running on the cluster.

In the following procedure, \$ANAHOME is the *pnuser* environment variable for the Prime Network installation directory (*/export/home/pnuser* by default). To install PN-IL on a server running Prime Network local redundancy software:

Step 1 On the primary cluster node (P1), log in as root and freeze the ana service.



Note The cluster server should be the active node where the ana service is running.

```
ssh root@active-cluster-node
clusvcadm -Z ana
```

Step 2 As *pnuser* (the operating system user for the Prime Network application), log into the active node where you froze the ana service.

```
su - pnuser
```

For example:

```
su - pn41
```

Step 3 Create an installation directory for PN-IL.

```
mkdir -p $ANAHOME/new-pnil-dir
```

For example, if the Prime Network installation directory was */export/home/pn41*, you would run this command to create an installation directory called *pnil*:

```
mkdir -p $ANAHOME/pnil
```

Step 4 Copy the installation files from the installation DVD, extract them, and start the installation script. These examples use the PN-IL installation directory named *pnil*.

a. Copy the PN-IL installation tar file from Disk 2 to the directory you created in [Step 3](#).

```
cp /tmp/sil-esb-2.2.0.tar.gz $ANAHOME/pnil
```

b. Change to the directory you created in [Step 3](#) and extract the PN-IL installation tar:

```
cd $ANAHOME/pnil
tar -zxf sil-esb-2.2.0.tar.gz
```


- c. Change to directory where the installation tar files were extracted and run the installation script:

```
cd sil-esb-2.2.0/install/packages
./installAndConfigureEsb.sh
```

- Step 5** Reload the user profile using the following command:

```
source $ANAHOME/.cshrc
```

Next, perform the necessary configuration steps that are described in [Configuring PN-IL on a Prime Network Gateway \(Local Redundancy\)](#), page 3-33.

Configuring PN-IL on a Prime Network Gateway (Local Redundancy)

If you are using Prime Network in standalone mode—that is, without Prime Central—configure PN-IL using the instructions in [Configuring PN-IL with Prime Network \(Standalone Mode with Local Redundancy\)](#), page 3-33.

If you are using Prime Network with Prime Central, configure PN-IL as described in [Configuring PN-IL with Prime Central \(Suite Mode with Local Redundancy\)](#), page 3-34.

Configuring PN-IL with Prime Network (Standalone Mode with Local Redundancy)

In standalone mode, Prime Network is not integrated with Prime Central and can independently expose MTOSI and 3GPP web services to other OSS/applications. In the following procedure, \$PRIMEHOME is the *pnuser* environment variable for the PN-IL installation directory you created in [Installing PN-IL on a Prime Network Server \(Local Redundancy\)](#), page 3-32.

- Step 1** As *pnuser*, configure PN-IL in standalone mode using the following command:

```
itgctl config 1 --anaPtpServer ana-cluster-ip --anaPtpUser pn-root-user --anaPtpPw
pn-root-user-password --authURL network-authentication-URL
```

itgctl uses these arguments.

Options and Arguments	Description
--anaPtpServer <i>ana-cluster-ip</i>	Specifies the virtual IP of the Prime Network primary cluster server.
--anaPtpUser <i>pn-root-user</i>	Specifies the name of Prime Network application root user (usually root).
--anaPtpPw <i>pn-root-user-password</i>	Specifies the password for Prime Network application root user.
--authURL <i>network-authentication-URL</i>	Specifies the URL used to authenticate Prime Network calls (usually https://localhost:6081/ana/services/userman).

For example:

```
itgctl config 1 --anaPtpServer 192.0.2.22 --anaPtpUser root --anaPtpPw myrootpassword
--authURL https://192.0.2.22:6081/ana/services/userman
```

Step 2 Start PN-IL by using the following command:

```
$PRIMEHOME/bin/itgctl start
```

Step 3 Log out as *pnuser* and log back in as the operating system root user.

Step 4 Unfreeze the ana service.

```
clusvcadm -U ana
```

Step 5 Enable NBI:

```
cd $PRIMEHOME/install/scripts  
./accessconfig.sh nbi enable
```

Next, disable the PN-IL health monitor as described in [Disabling the PN-IL Health Monitor, page 3-36](#).

Configuring PN-IL with Prime Central (Suite Mode with Local Redundancy)




Note Use this procedure only after installing the PN-IL as described in [Installing and Configuring PN-IL with Local Redundancy, page 3-30](#).

When Prime Network is in suite mode, it is integrated with Prime Central. This procedure explains how to integrate PN-IL with a deployment of Prime Central that is using gateway local redundancy. You can use this procedure for:

- New installations of PN-IL with local redundancy.
- Existing standalone installations of PN-IL with local redundancy, that you want to move from standalone to suite mode.

[Figure 3-2](#) illustrates the deployment of local redundancy in Suite Mode.

Figure 3-2 Local Redundancy Suite Mode



In the following procedure, \$PRIMEHOME is the *pnuser* environment variable for the PN-IL installation directory you created in [Installing PN-IL on a Prime Network Server \(Local Redundancy\)](#), page 3-32.

Before You Begin

Make sure Prime Network is in suite mode. For information on integrating Prime Network with Prime Central, refer to the [Cisco Prime Central Quick Start Guide](#).

Step 1 Edit the necessary integration files and run the integration script:

- a. Log into the Prime Network primary gateway server as *pnuser* and change to the \$PRIMEHOME/integration directory.

```
cd $PRIMEHOME/integration
```

- b. Edit the **ILIntegrator.prop** file and change the value of the 'HOSTNAME' property to ana-cluster-ana, which is the fixed name for the Prime Network cluster server.

```
HOSTNAME=ana-cluster-ana
```

- c. Execute the following integration script to integrate PN-IL into the deployment:



Note When you run DMIntegrator.sh, you must exactly follow the format below or the script will fail.

```
./DMIntegrator.sh -a ILIntegrator.prop prime-central-db-hostname
prime-central-db-service-name prime-central-db-user prime-central-db-user-password
prime-central-port-number
```

DMIntegrators.sh uses these variables. You must enter them in this exact order.

DMIntegrator.sh Variable	Description
<i>prime-central-db-hostname</i>	Specifies the IP address of the Prime Central database server
<i>prime-central-db-service-name</i>	Specifies the name of Prime Central database service
<i>prime-central-db-user</i>	Specifies the name of Prime Central database user (usually primedba)
<i>prime-central-db-user-password</i>	Specifies the password for Prime Central database user
<i>prime-central-db-port</i>	Specifies the port for Prime Central database (usually 1521)

Example:

```
./DMIntegrator.sh -a ILIntegrator.prop 10.10.10.10 primedb primedba mypassword 1521
```

Step 2 Reload the user profile:

```
source $PRIMEHOME/.cshrc
```

Step 3 Start PN-IL:

```
$PRIMEHOME/bin/itgctl start
```

Step 4 Log out as pnuser and log back in as the operating system root user.

Step 5 Unfreeze the ana service.

```
clusvcadm -U ana
```

Step 6 Enable NBI:

```
cd $PRIMEHOME/install/scripts
./accessconfig.sh nbi enable
```

Next, disable the PN-IL health monitor as described in [Disabling the PN-IL Health Monitor, page 3-36](#).

Disabling the PN-IL Health Monitor

When PN-IL is installed in a local redundancy deployment, the RHCS cluster service monitors PN-IL's status. Therefore, you should disable the PN-IL health monitor.

To disable the PN-IL health monitor, log in as *pnuser* and execute the following command:

```
$PRIMEHOME/local/scripts/il-watch-dog.sh disable
```