CHAPTER 1

Gateway High Availability Overview

Prime Network provides gateway redundancy for local and geographical sites. The Prime Network high availability solutions uses Red Hat Cluster Suite (RHCS) for local and Oracle Active Data Guard (ADG) for geographical redundant configurations. The gateway high availability solution is available in following configurations:

- Local redundancy only
- Geographical redundancy only
- Local and Geographical redundancy

The local redundancy configuration uses two active local servers to provide an automatic failover solution and the geographical redundancy configuration adds an additional server at a geographical site for a full disaster recovery solution.

Note If Prime Network is installed in a local redundancy, RHCS HA configuration, you can integrate it with Prime Central. However, if Prime Network is installed in a geographic redundancy HA configuration, you cannot integrate it with Prime Central. For steps on integrating Prime Network HA with Prime Central, see *Cisco Prime Network 1.1 Quick Start Guide.*

For gateway server high availability solution that uses Veritas software, contact your Cisco account representative for further assistance.

Scope of the Guide

This section describes the assumptions upon which the information in this guide is based. If your high availability deployment differs from what is described here, please contact your Cisco account representative for assistance with planning and installation of high availability.

The high availability solutions discussed in this guide are based on the following assumptions:

- Your setup uses the Prime Network embedded database (Oracle 11g R2), not an external database.
- Does not support IPv6 on the gateway or database.

- Does not include HA protection to the Prime Network units. For details on unit server high availability, see the *Cisco Prime Network 3.10 Administrator Guide*.

  Both the local redundancy and geographic redundancy configurations are independent of and compatible with the unit server high availability mechanism described in the administrator guide.

- Does not include the upgrade procedure. For upgrade procedure, see *Cisco Prime Network 3.10 Installation Guide.*

Local Redundancy Functional Overview

The local redundancy configuration uses one active and one standby node to provide an automatic failover solution for local hardware faults without the need to reconfigure IP addresses. The solution uses the Red Hat Cluster Suite (RHCS). The nodes are monitored by RHCS and if the node managing the services fails, the services are seamlessly moved to the other node. 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.

When this solution is initially installed, the gateway and database services are installed on and managed by one node in the cluster from where the installation script is run.

Figure 1-1 shows a basic dual-node cluster local redundancy configuration, where the Prime Network gateway service is on Server P1, the Oracle database service is on Server P2, and both servers are connected to an embedded database that is installed on an external device.

The RHCS local redundancy solution requires a fencing device, which is a hardware unit that disconnects a node from shared storage to ensure data integrity. For more information on fencing options, see Fencing Option, page 1-4.

Figure 1-1 Architecture for Gateway with RHCS Local Redundancy

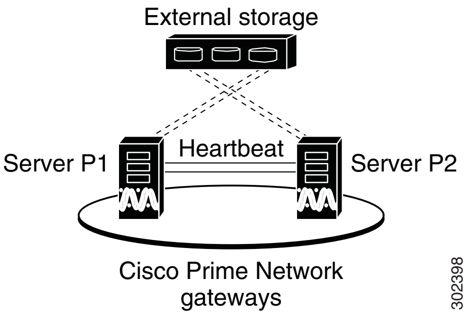

Configuration Details for Local Redundancy

Local redundancy requires that RHCS be installed on both nodes. Out of the box, both services run on the node from which the installation script is run. This configuration can be changed, if desired, using RHCS web GUI or CLI (clusvcadm utility). For details on the required system configuration for local redundancy, see

The local redundancy setup has the following:

- Dual Node Cluster, page 1-3
- RHCS Installed on Both Nodes, page 1-3
- External Shared Storage, page 1-3
- Fencing Option, page 1-4
- Security, page 1-5

Dual Node Cluster

The Prime Network gateway and embedded database are installed in a dual-node cluster. Each node has the platform to run both Prime Network gateway and database services.

RHCS Installed on Both Nodes

RHCS manages the local redundancy by monitoring cluster configured services: **ana** and **oracle_db**. If a hardware or software failure occurs, the RHCS automatically restarts the failed node's services on the functional node.

Table 1-1 lists the services that are monitored by RHCS.

Table 1-1 Cluster Configured Services Monitored by RHCS

RHCS Service	Description	
ana	Monitors AVM 99 (Prime Network) and consists of the following resources.	
	IP address	*ana_service_floating_IP*
	Scripts	ana.sh
oracle_db	Monitors Oracle processes and listener and consists of the following resources.	
	IP address	*oracle_db_floating_IP*
	Scripts	oracles.sh, lsnr.sh

The Oracle listener should be running before Prime Network, which allows the Prime Network gateway process (AVM 11) to connect to the database. If the listener is not running, the Prime Network agent contains logic to enable it to delay startup of the Prime Network processes while it waits for the listener to start. If the listener does not start up on time, the Prime Network gateway agent will abort the startup, resulting in a Prime Network resource failure.

Alternatively, you can also bring the service groups online in serial sequence, starting with the Oracle service group, then the Prime Network service group. (RHCS does not enforce this behavior.)

External Shared Storage

RHCS requires an external shared storage that is mountable from both nodes. The external shared storage contain the Prime Network files and the Oracle files.

Fencing Option

Each node in the cluster must use a fencing method. Local redundancy configuration uses a fencing hardware unit for cutting a node off from the shared storage. This is to ensure data integrity and to prevents a “split brain” scenario by preventing the problematic node from writing to the shared storage. If any problem with cluster node occurs, RHCS invokes the fencing device with the peer and waits for the success signal. If a failure occurs, the cut off can be accomplished by powering off the node with a remote power switch, disabling a switch channel, or revoking a host's SCSI 3 reservations.

The supported fencing options are:

- fence_ipmilan— Intelligent Platform Management Interface (IPMI) v1.5 or higher compliant devices over a LAN.
- fence_ilo— Hewlett Packard Integrated Lights Out (HP iLO).
- fence_manual—This option allows you to add a a Red Hat-supported fencing device not listed above. If you choose Manual, the fence-manual-fencing agent is assigned. This fencing agent is temporary and should not be used in production because it does not perform automated actions. If a cluster problem occurs, the node and storage must be manually disconnected, or another fencing agent must be used to disconnect them. If you choose this option during the installation because you want 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, as shown in Figure 1-2.

Note General information about the RHCS web GUI is provided in Configuring the RHCS Web Interface (Optional), page 2-16. However, see the Red Hat *Conga User Guide* for complete information about using the RHCS web GUI application. Additionally, you need the RHCS user documentation to provision and manage cluster fencing devices.

Figure 1-2 RHCS GUI Fencing Method Window

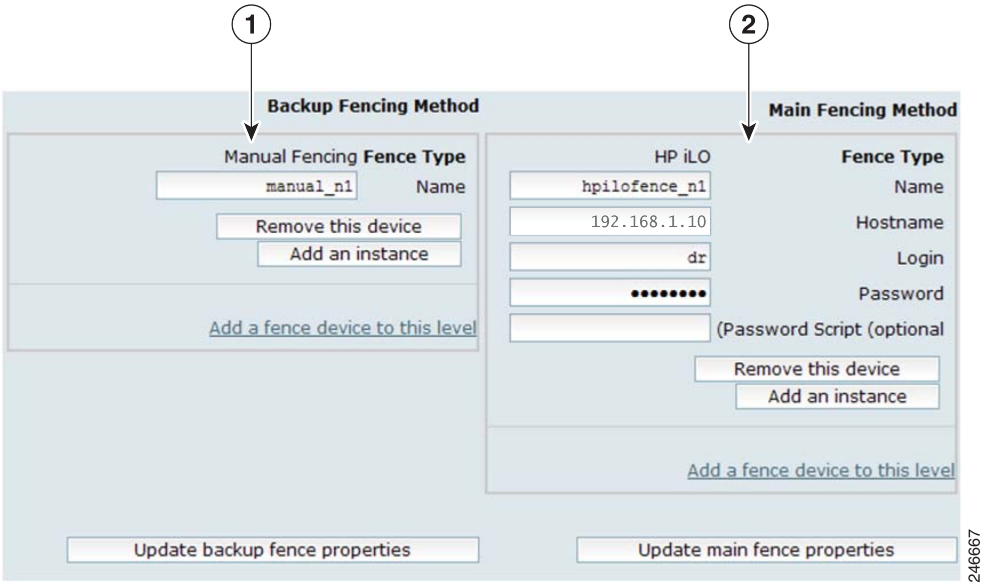

Backup Fencing Method

Manual Fencing Fence Type
manual_n1 Name

Main Fencing Method

HP iLO Fence Type
hpilofence_n1 Name
192.168.1.10 Hostname
dr Login
•••••••• Password
(Password Script (optional

1
2
Remove this device
Add an instance
Add a fence device to this level
Remove this device
Add an instance
Add a fence device to this level
Update backup fence properties
Update main fence properties
246667

1	Backup fencing method: Move the manual fencing agent to the backup method.	**2**	Main fencing method: Add a different Red Hat-supported fencing device.

Note To prevent fencing loops, the cluster interconnect and power fencing (for example, HP-iLO) should use the same network, such as bond0.

Note If the main fencing device is a remote power switch, define all ports that supply power to the node simultaneously.

Security

When the RHCS local redundancy solution is installed, SSL keys are generated and copied to the other node in the cluster.

Geographical Redundancy Functional Overview

The geographical redundancy of the Prime Network gateway is implemented using Oracle Active Data Guard (ADG). The ADG geographical redundancy solution uses a remote site containing a single server that provides failover in case of a failure at the primary site. The remote site, which is running but has no active applications, provides redundancy for the server (or servers) at the primary site, which contain the gateway and the database services.

Note In case of both local and geographical redundancy, there will be a dual-node cluster in the primary site with external shared storage, whereas in case of only geographical redundancy, the primary site will have a single node with external storage. This section covers the scenario when dual-node is deployed at the primary site.

Figure 1-3 shows the dual-node, local redundant cluster with a remote single-node site, where:

- Primary site (P1 and P2) is a dual-node cluster with an external shared storage.
- Remote site (S1) is a single-node with an external storage. Server S1 contain its own server, database, and storage, all located at another geographical location. The remote site will be the backup to the first site.

Figure 1-3 Geographical Redundancy Setup

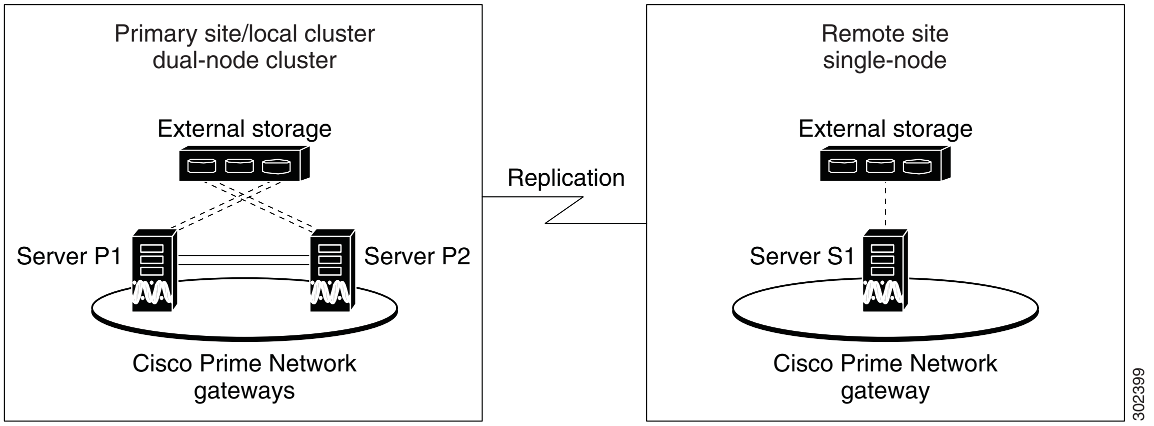

Note Geographical redundancy does not allow the Prime Network service (ana) to be brought online on the local side while the Oracle service is online on the remote site (or vice versa).

The data stored in the server and database is continuously replicated between the two sites. The primary and standby database are monitored and synchronized using ADG; the Prime Network server files (registry and system files) are synchronized using the GWSync utility, which is based on Red Hat Enterprise Linux rsync. Prime Network periodically monitors and validates the replication process and issues a System event in case of a problem. Figure 1-4 shows the data replication process between the primary site and remote site. To secure the channel used for data replication, an SSH key exchange is performed during the Prime Network installation.

Figure 1-4 Hardware Configuration for Geographical Redundancy

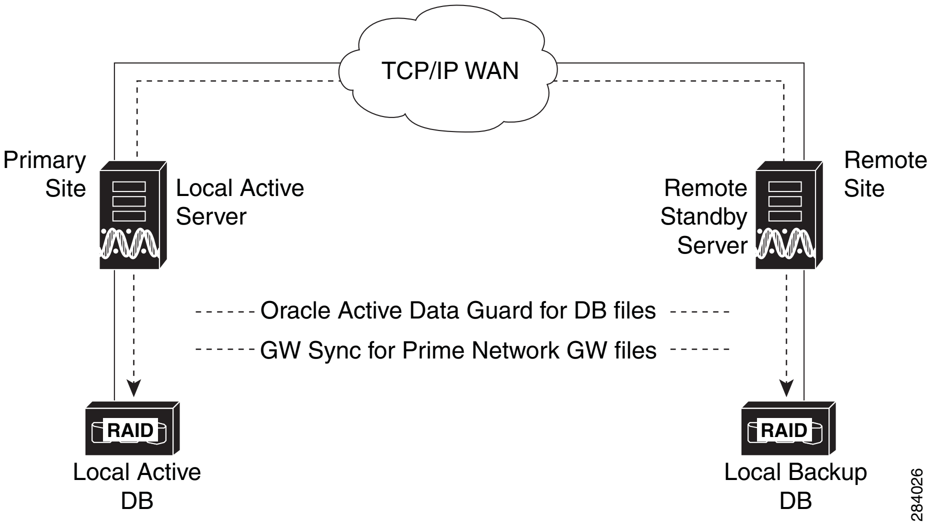

For disaster recovery (if the primary site becomes unavailable), a manual failover can be triggered from the remote site. The utilities for managing the manual failover are described in Maintaining Geographical Redundancy, page 3-14.

The geographical redundancy solution uses two replication processes.

- Oracle ADG Replication Process, page 1-6
- GWSync Replication Process, page 1-7

Oracle ADG Replication Process

When the ADG solution is installed, a standby database is created at the remote site to replicate Prime Network database information. The remote site database is an active (read-only) Oracle instance. The primary site database, which operates in archive log mode, sends copies of the redo logs to the remote site database for archiving. Data is synchronized using Redo-apply.

When the high availability solution is installed, it sets up the cron jobs that will monitor the synchronization process. ADG uses port 1521 for communication between the servers. This port must be open.

Figure 1-5 illustrates how data is replicated between the primary site database and the remote site.

Figure 1-5 ADG Database Replication Process (ADG Geographical Redundancy)

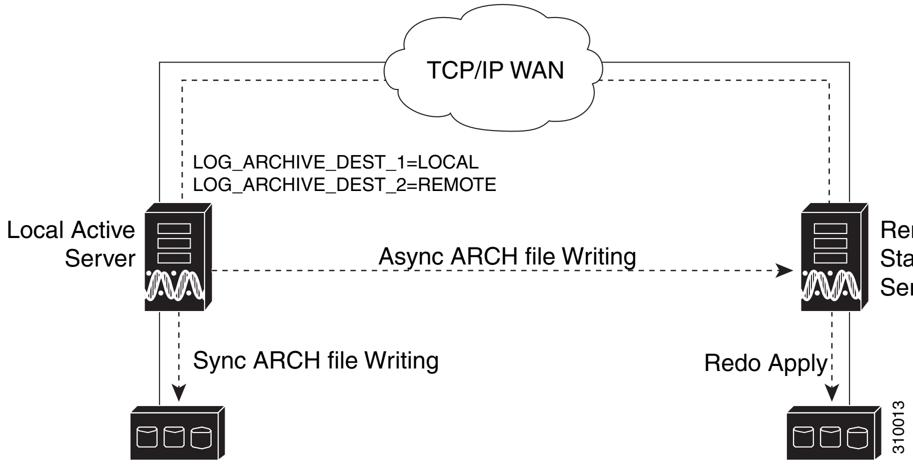

Note The databases must have identical disk capacities and mount points.

To troubleshoot problems with the replication process, see the Verifying the Geographical Redundancy Setup, page 3-10.

GWSync Replication Process

Gateway Sync is a RHEL rsync utility that replicates the Prime Network home directory (and any file system data that is required for disaster recovery) from the primary gateway to the remote site. Cron jobs trigger synchronization at both the primary and remote sites. Data is exchanged using SSH across secure channels.

Data is sent on an incremental basis. In other words, GWSync only sends data that has changed.

The initial GWSync is triggered when the geographical redundancy solution is installed; after that, the data is synchronized every 60 seconds. The installation process also sets up the cron jobs that trigger the synchronization process

Prime Network monitors and validates the replication processes (ADG and Gateway Sync) and issues a system event if replication problems occur. To troubleshoot problems with the replication process, see Verifying the Geographical Redundancy Setup, page 3-10.

Figure 1-6 How GWSync Replication Process is Monitored (ADG Geographical Redundancy)

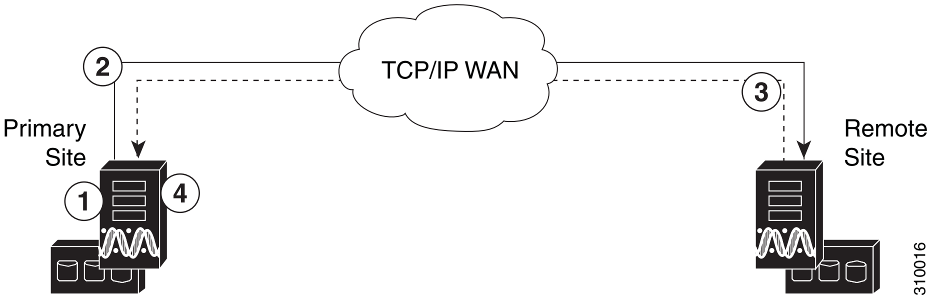

1	Local primary site generates local_timestamp file.	**3**	Primary site pulls remote site's timestamp file as remote_timestamp.
2	remote site pulls *NETWORKHOME* directory from local primary site (including remote site's local_timestamp file).	**4**	Primary site compares local_timestamp and remote_timestamp files and, if too much time has passed, issues a System event.

Licenses

A Prime Network base license must be registered and activated within 120 days of installation. In case of upgrade, use the existing license files as new license files will not be generated when upgrading from one version to the other.

If there are no valid license files on the gateway server, Prime Network acts as an evaluation version. This means that it has full functionality for 120 days after installation, and then it expires and UI connections will be disallowed.

The license directory *NETWORKHOME*/Main/ha/licenses on the active gateway should contain a copy of all license files for all servers. This directory will be available to all servers because it is part of the Prime Network partition that is replicated among servers. If you add new licenses, you must copy them to this directory and run the **resetLicenses.pl** command to read the licenses.

Do the following if you are using gateway server high availability:

Step 1 For new installation, ensure you have the required stand-by part number to use the Prime Network product in HA mode. If you are upgrading from a previous version, continue to the next step.

Note If you do not have the stand-by part number, contact your Cisco account manager and ask them to provide the license including licenses for the additional gateway server(s) if there are any.

Step 2 Place a copy of all gateway server license files in the license directory *NETWORKHOME*/Main/ha/licenses on an active gateway. This directory is part of the Prime Network partition that is shared (local redundancy) or replicated (geographical redundancy) between servers in the gateway server high availability solution. If you add new licenses, you must copy them to this directory on the active gateway.

Note All gateway licenses should be copied to *NETWORKHOME*/Main/ha/licenses on the active gateway.

Step 3 After installing new license files, on the active gateway enter the following commands to make sure that the files are detected (You only need to do this once)

```
liccontrol stop
perl $ANAHOME/Main/ha/resertLicenses.pl
liccontrol start
```

Gateway High Availability Installation Scripts

This section includes the details of the installation scripts (**install_Prime_HA.pl** and **setup_Prime_DR.pl**) used for configuring local and geographical redundancy. These scripts are provided in the **RH_ha.zip** file on the Disk 1: New Install DVD.

In Table 1-2, x mark indicates the script should be run only on the marked server and N/A indicates not applicable for that server.

Table 1-2 High Availability Installation Scripts

Scripts	Description	Primary Site		Remote Site
		Server P1 **(Primary database)**	**Server P2**	**Server S1**
install_Prime_HA.pl	Installs local redundancy elements. • If local redundancy is configured, then this script is triggered from a locally mounted node. • Installs gateway and embedded database on the locally mounted node. • Partial installation of local redundant node element including cron jobs, users, groups, SSH keys and other elements. • If there is a remote site, this script will install the Prime Network gateway and embedded database on a remote site.	x	N/A	N/A
setup_Prime_DR.pl	Configures the geographical redundant server. • Stops the Cisco Prime Network gateway for the configuration process. • Configures the Oracle ADG. • Sets up cron jobs for Gateway Sync. Sets up cron jobs for ADG monitoring. • Triggers the initial Gateway Sync. **Note** The **setup_Prime_DR.pl** script must run on the node running the primary database.	x	N/A	N/A

Gateway High Availability Installation Flow

Table 1-3 and Figure 1-7 provides the general flow in which to complete the setting up Prime Network 3.10 high availability. For detailed instructions on each, see the topics referenced in Table 1-3.

Table 1-3 Prime Network Gateway High Availability Installation Flow

	Procedure	Flow
Step 1	Determine the type of HA installation. See Local Redundancy Functional Overview, page 1-2 and Geographical Redundancy Functional Overview, page 1-5.	Proceed to next step.
Step 2	Do the gateway where you will install Prime Network 3.10 High Availability meet the local or geography redundancy requirements specified in Prerequisites for Local Redundancy, page 2-1 and Prerequisites for Geographical Redundancy, page 3-1?	**Yes:** Go to next step. **No:** Do not continue until all specified requirements are met.
Step 3	Have you completed all preinstallation tasks in Setting Up Local Redundancy, page 2-3 and section Setting Up Geographical Redundancy, page 3-3?	**Yes:** Go to next step. **No:** Do not continue until all preinstallation tasks are completed.
Step 4	Have you completed setting up HA on gateway for local and remote site? See Installing the Gateway for Local Redundancy, page 2-11and Installing the Gateway for Geographical Redundancy, page 3-4.	**Yes:** Go to next step. **No:** Do not continue until all setting up instructions are completed.
Step 5	Have you completed verifying local and geographical setup? See Verifying the Local Redundancy Setup, page 2-19 and Verifying the Geographical Redundancy Setup, page 3-10 section.	**Yes:** Go to next step. **No:** Complete verifying the setup.
Step 6	Installation completed.	

Flow Diagram

Figure 1-7 Local and Geographical Redundancy Flow

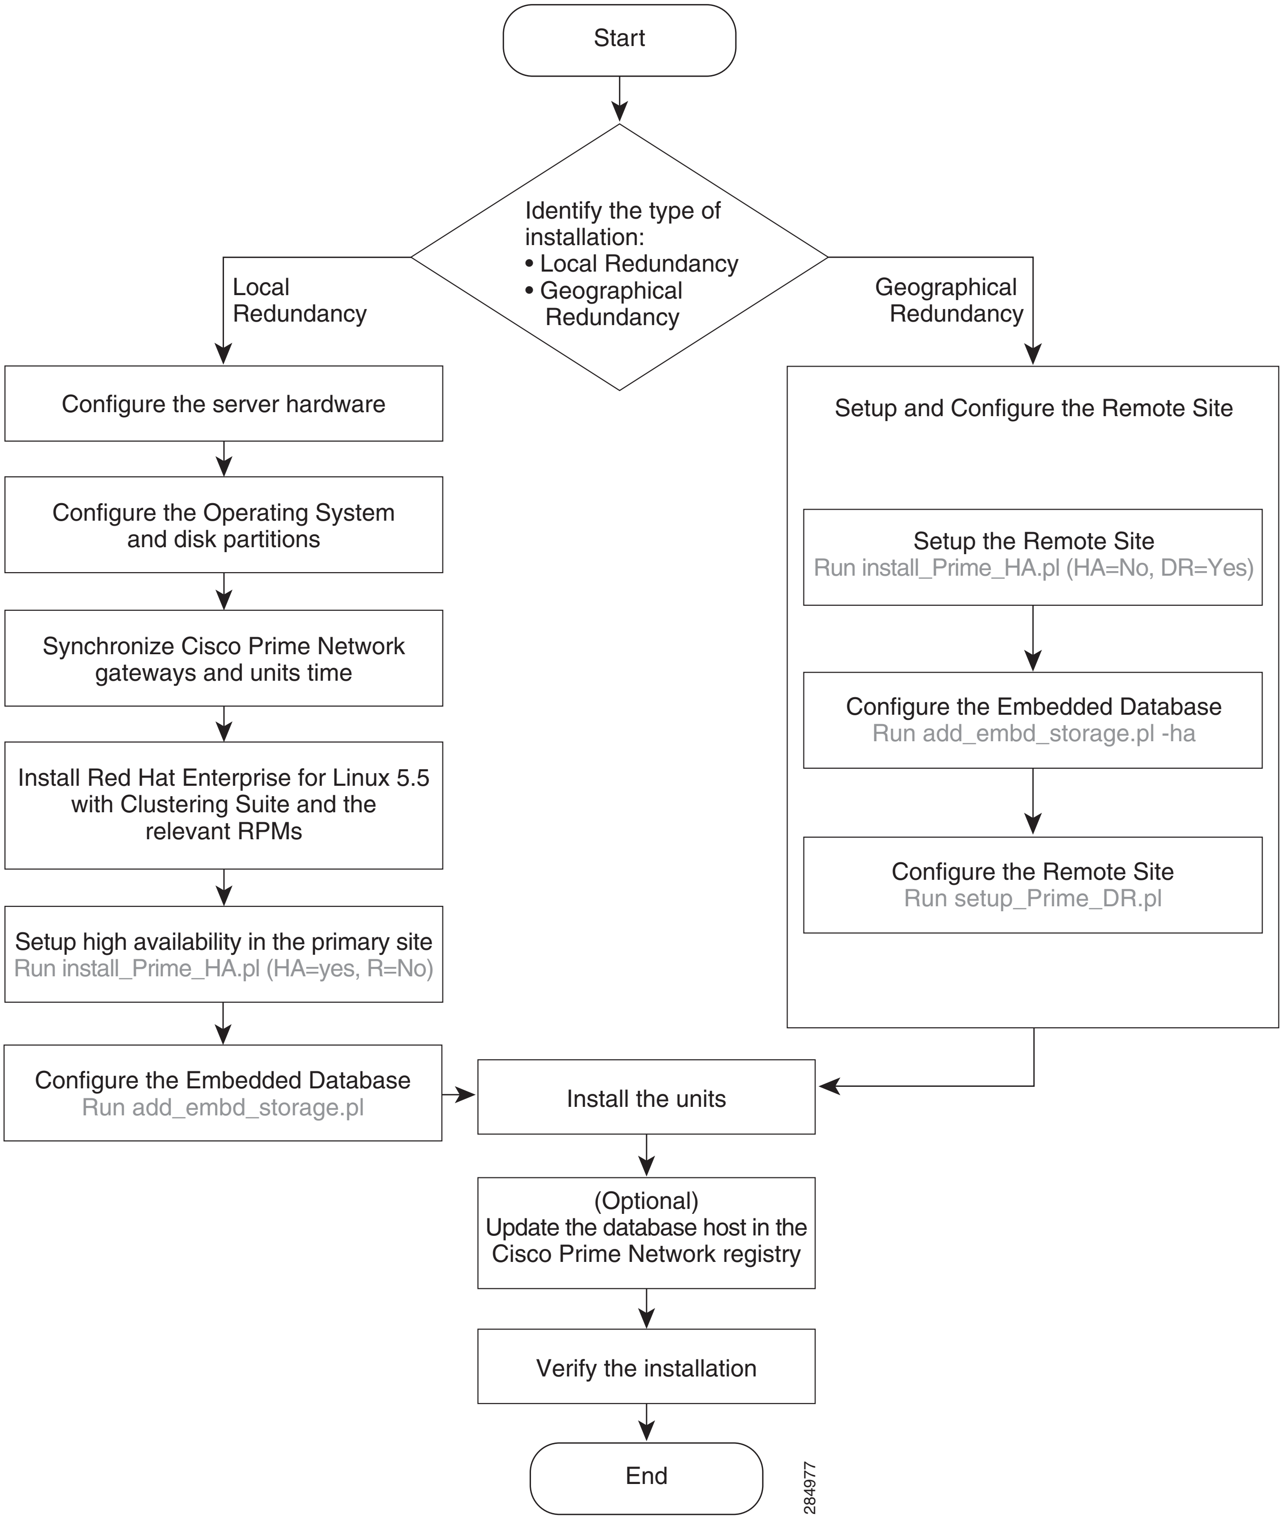
Start
Identify the type of installation:
• Local Redundancy
• Geographical Redundancy
Local Redundancy
Geographical Redundancy
Configure the server hardware
Configure the Operating System and disk partitions
Synchronize Cisco Prime Network gateways and units time
Install Red Hat Enterprise for Linux 5.5 with Clustering Suite and the relevant RPMs
Setup high availability in the primary site
Run install_Prime_HA.pl (HA=yes, R=No)
Configure the Embedded Database
Run add_embd_storage.pl
Setup and Configure the Remote Site
Setup the Remote Site
Run install_Prime_HA.pl (HA=No, DR=Yes)
Configure the Embedded Database
Run add_embd_storage.pl -ha
Configure the Remote Site
Run setup_Prime_DR.pl
Install the units
(Optional)
Update the database host in the Cisco Prime Network registry
Verify the installation
End
284977