CHAPTER 18

Using RHCS/ADG Gateway Server High Availability

These topics describe the gateway server high availability solutions that use Red Hat Cluster Suite (RHCS) and Oracle Active Data Guard (ADG). These solutions leverage the Prime Network embedded database and existing licensing terms. Use the architecture described in these topics as a reference point and adjust them to meet the needs of specific deployments. Both the local redundancy and geographic redundancy configurations are independent of and compatible with the unit server high availability mechanism (described in Unit Server High Availability and AVM Protection, page 16-1).

Note This solution does not support IPv6 on the gateway or database. This solution also does not support a remote database. (In other words, for local redundancy, the database must be installed on the same server as the gateway. For geographical redundancy, the database must be installed on the standby server (with the gateway)).

For information on the gateway server high availability solution that uses Veritas software, see Using Veritas Gateway Server High Availability, page 17-1.

Red Hat Cluster Suite (RHCS) Local Redundancy

The RHCS local redundancy solution contains a dual-node cluster that provides an automatic failover solution for local hardware faults. Because the gateway and database use logical IP addresses (which they retain regardless of the node they are running on), if a failover occurs, there is no need to reconfigure IP addresses.

Note This solution does not support IPv6 on the gateway or database. This solution also does not support a remote database. (In other words, for local redundancy, the database must be installed on the same server as the gateway. For geographical redundancy, the database must be installed on the standby server (with the gateway)).

When this solution is initially installed, the gateway and database services are installed on and managed by one node in the cluster. The nodes are monitored by RHCS and if the node managing the services fails, the services are seamlessly moved to the other node.

If desired, one of the services can be moved to the other node using the RHCS web GUI or CLI (**clusvcadm** utility). This type of configuration is shown in Figure 18-1, 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.

Note You cannot use the local redundancy configuration that is illustrated in Figure 18-1 for geographical redundancy because geographical redundancy requires a WAN and a dedicated connection from the gateway to the database.

Figure 18-1 Architecture for Gateway with RHCS Local Redundancy

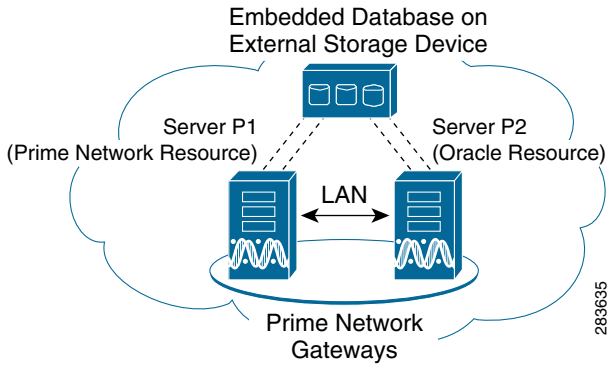

The RHCS local redundancy solution also requires a *fencing* device, which is a hardware unit that disconnects a node from shared storage to ensure data integrity. For information on the supported fencing options, see the *Cisco Prime Network 3.8 Installation Guide*.

To troubleshoot problems with hardware and service failures, see the *Cisco Prime Network 3.8 Installation Guide*.

Configuration Details for RHCS Local Redundancy

Local redundancy requires that Red Hat Cluster Suite (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).

The local redundancy solution also requires a fencing hardware unit for cutting a node off from the shared storage. This ensures data integrity and prevents a split brain scenario, where the node are disconnected from each other and each presumes the other has failed. 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.

Note For complete redundancy, a configuration with no single point of failure is recommended.

In cases where your configuration has more than one port connection, you should consider adding a redundant fencing device. To troubleshoot problems with hardware failures, see the *Cisco Prime Network 3.8 Installation Guide*.

Fencing Devices

Each node in the cluster must use a fencing method. The fencing method is engaged when one of the nodes has a problem, and prevents the problematic node from writing to the shared storage. A complete list of supported fencing devices is provided the *Cisco Prime Network 3.8 Installation Guide*. Manual fencing is also supported but is recommended as a temporary solution; for more information, see Fencing and Manual Fencing, page 18-7.

RHCS Services and Resources

Services are a set of resources that are grouped together. RHCS monitors two services: **ana** and **oracle_db**.

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.)

Table 18-1 lists the services that are monitored by RHCS.

Table 18-1 Local Redundancy 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

Both nodes in the cluster must have identical RHCS versions and packages.

For information on the RHCS version that is supported and to troubleshoot service failures, see the *Cisco Prime Network 3.8 Installation Guide*.

IP Addresses

Both the Prime Network gateway process and the Oracle database services have their own virtual IP address (floating IPs). Because they retain their IP addresses when there is a failover or switchover, Prime Network clients interprets failovers or switchover as local service restarts. Virtual IP addresses must be in the same subnet. Use the gateway floating IP address when installing a new unit, using LDAP, etc.

Note If you are using the **network-conf** script, when you are prompted for the IP address of units, use the floating IP address of the gateway.

Multicast Addresses

Each network switch and associated networking equipment in a Red Hat cluster must support and enable multicast addresses and IGMP. Because configuration procedures vary, refer to the appropriate vendor documentation or other information about configuring network switches, and associated networking equipment, to enable multicast addresses and IGMP.

Because the Prime Network installer does not validate the multicast address, you must manually verify that the multicast address meets RHCS requirements and is not blocked by a firewall.

File System Type

The RHCS local redundancy solution requires the ext3 file system.

Shared Storage and Disk Partitioning

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.

The storage related to the services managed by the cluster should not be automatically mounted (automounted) by the operating system upon reboot. RHCS will perform the mounting.

Security

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

Licenses

The license directory *NETWORKHOME*/Main/ha/licenses on the active gateway should contain a copy of all license files for both nodes. This directory will be available to both nodes because it is part of the partition that is shared. If you add new licenses, you must copy them to this directory and run the **resetLicenses.pl** command to read the licenses. See Licensing and Gateway Server High Availability, page 5-2.

Failover and Switchover (RHCS 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.

Human intervention is required only in exceptional cases, such as when the database becomes corrupted or a component fails, and the component is not configured for redundancy. Manual switchovers are performed using the RHCS web GUI or the **clusvcadm** utility. Once a failed node is repaired, you must perform a manual switchover to revert the cluster to its original configuration. See Managing RHCS Local Redundancy, page 18-5

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

Managing RHCS Local Redundancy

Note Before stopping the Prime Network or Oracle application processes, place the RHCS services in maintenance mode (also known as *freezing* the process) using **clusvcadm**. If you attempt to restart either the Prime Network or Oracle applications without freezing the RHCS process, the cluster may detect that the services are down and attempt to restart them. See Stopping and Restarting RHCS Services Using the RHCS CLI (clustat, clusvcadm), page 18-6.

These topics provide information pertaining to ongoing management of an RHCS local redundancy cluster.

RHCS 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 csi-bvc clurgmgrd: [27961]: <info> Executing /usr/local/bin/ana.sh status
Mar 23 13:46:07 csi-bvc clurgmgrd: [27961]: <info> Executing /usr/local/bin/oracle.sh
status
Mar 23 13:46:07 csi-bvc clurgmgrd: [27961]: <info> Executing /usr/local/bin/lsnr.sh status
```

To troubleshoot problems with service failures, see the *Cisco Prime Network 3.8 Installation Guide*.

RHCS Web GUI (luci)

The RHCS web client provides information about the status of the cluster (the status of each service, the node the service is running on, and so forth). You can also use the web GUI 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).

Note To stop or start the Prime Network and Oracle database services that are managed by the cluster, use **clusvcadm**, as described in Stopping and Restarting RHCS Services Using the RHCS CLI (clustat, clusvcadm), page 18-6.

The RHCS web interface is automatically configured by the Prime Network installation script. 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**

For details on how to use the web GUI, see the appropriate RHCS documentation.

Stopping and Restarting RHCS Services Using the RHCS CLI (clustat, clusvcadm)

Note Before stopping the Prime Network or Oracle application processes, place the RHCS services in maintenance mode (also known as *freezing* the process) using **clusvcadm**. If you attempt to restart either the Prime Network or Oracle applications without freezing the RHCS process, the cluster may detect that the services are down and attempt to restart them.

The **clustat** and **clusvcadm** commands are the basic CLI commands you can use to monitor and manage the local redundancy cluster. This topic describes some common uses for these commands. You must be logged in as root to use these commands.

The **clustat** command checks a cluster's members and overall status. In the following example, the cluster name is ana_cluster and csi-bvc.cisco.com is the node from which the command was run.

```
root@csi-bvc.cisco.com]# clustat
Cluster Status for ana_cluster @ Thu Mar  3 10:24:50 2011
Member Status: Quorate

 Member Name                                              ID   Status
 ------ ----                                              --   ------
 csi-bvc.cisco.com                                          1   Online, Local, rgmanager
 csi-w47.cisco.com                                          2   Online, rgmanager

 Service Name                                    Owner (Last)                    State
 ------- ----                                    ----- ------                    -----
 service:ana                                     csi-bvc.cisco.com               started
 service:oracle_db                               csi-w47.cisco.com               started
```

If you need to restart Prime Network or the Oracle application processes, first use the **clusvcadm** command to stop the RHCS services using 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** or **oracle_db**.

```
# 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@csi-bvc.cisco.com]# clustat
Cluster Status for ana_cluster @ Thu Mar  3 12:31:55 2011
Member Status: Quorate

 Member Name                                              ID   Status
 ------ ----                                              --   ------
 csi-w47.cisco.com                                          1   Online, rgmanager
 csi-w47.cisco.com                                          2   Online, Local, rgmanager

 Service Name                                    Owner (Last)                    State
 ------- ----                                    ----- ------                    -----
 service:ana                                     csi-w47.cisco.com               started [Z]
 service:oracle_db                               csi-w47.cisco.com               started [Z]
```

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

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

```
# clusvcadm -U service
```

Fencing and Manual Fencing

A fencing device is a hardware unit that disconnects a node from the shared storage. This happens when a node needs to assume control of a service but cannot connect to the other node. Disconnecting the problematic node from the database ensures data integrity and prevents split-brain scenarios. You can reconfigure the fencing choice at any time using the RHCS web interface or other RHCS tools.

During the installation of the RHCS solution, you are prompted to select one of four fencing options. The first three are for specific fencing devices supported by the solution. If you choose one of these devices (or more specifically, one of these fencing agents), if an error occurs, the fencing agent will automatically disconnect the cluster node from the storage.

The fourth option is *manual fencing*. If you choose manual fencing, this means you are responsible for making sure that, when a problem occurs, the node and storage are disconnected (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 you use manual fencing 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. When you run it, this command asks for confirmation that you have disconnected the faulty node from the storage. Only then will the cluster workflow continue.

```
Warning:  If the node "csi-bvc.cisco.com" has not been manually fenced
(i.e. power cycled or disconnected from shared storage devices)
the GFS file system may become corrupted and all its data
unrecoverable!  Please verify that the node shown above has
been reset or disconnected from storage.

Are you certain you want to continue? [yN] y
```

To use the **fence_ack_manual** command, log into the gateway server as root and enter the command using the following syntax. The node that has been disconnected from storage is specified using the **-n** *nodename* option.

fence_ack_manual -n *nodename*

To troubleshoot hardware failures, see the *Cisco Prime Network 3.8 Installation Guide*.

Oracle Active Data Guard (ADG) Geographical Redundancy

The ADG geographical redundancy solution uses a secondary site containing a single server that provides failover in case of a failure at the primary site. The remote secondary server, 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.

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 Oracle Active Data Guard; 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.

Note This solution does not support IPv6 on the gateway or database. This solution also does not support a remote database. (In other words, for local redundancy, the database must be installed on the same server as the gateway. For geographical redundancy, the database must be installed on the standby server (with the gateway)).

For disaster recovery (if the primary site becomes unavailable), a manual failover can be triggered from the standby site.

The gateway and database use logical IP addresses which are different between the two sites (the sites are most likely on different subnets). The utilities for managing the manual failover are described in Managing ADG Geographical Redundancy, page 18-14.

Figure 18-2 illustrates a geographical redundancy with the following members:

- A primary site, with Server P1 containing the Prime Network gateway service, and Server P2 containing the Oracle database service. Both servers are connected to an embedded database that is installed on an external device.
- A remote site, with Server S1 containing its own server, database, and storage, all located at another geographical location. The secondary site will be the backup to the first site.

Figure 18-2 Architecture for Gateway with ADG Geographic Redundancy

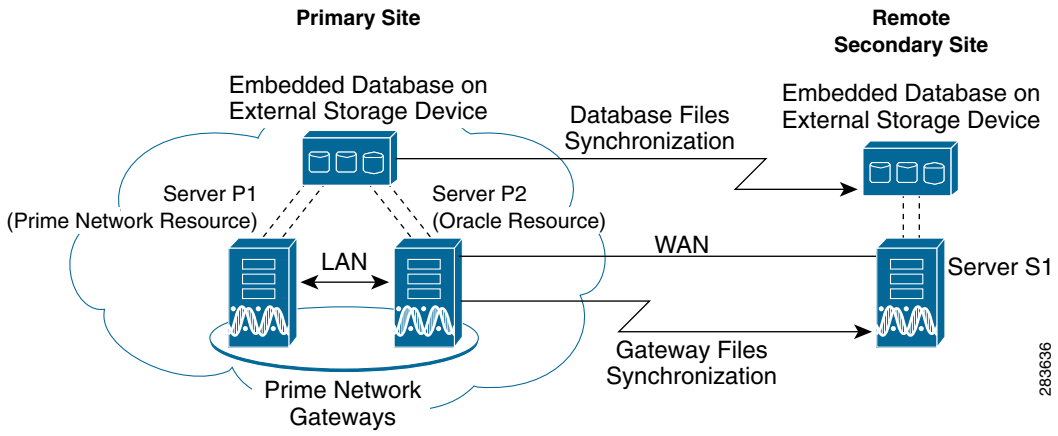

Configuration Details for ADG Geographical Redundancy

The ADG geographical redundancy solution requires the following applications:

- Oracle Active Data Guard Option between the primary local and secondary remote machine—Replicates data to a standby database at remote site (the standby database is set up during installation of the ADG solution). See Oracle ADG Replication Process and Configuration Files, page 18-9.
- GWSync—Replicates the server home directory (and any file system data that is required for disaster recovery) to the server at the remote site. See GWSync Replication Process, page 18-11.

In addition, you must enable backups for the embedded database, as described in Enabling Backups (Embedded Database), page 11-11.

You cannot use the local redundancy configuration that is illustrated in Figure 18-1 on page 18-2 for geographical redundancy because geographical redundancy requires a WAN and a dedicated connection from the gateway to the database.

Figure 18-3 Hardware Configuration for ADG Geographical Redundancy

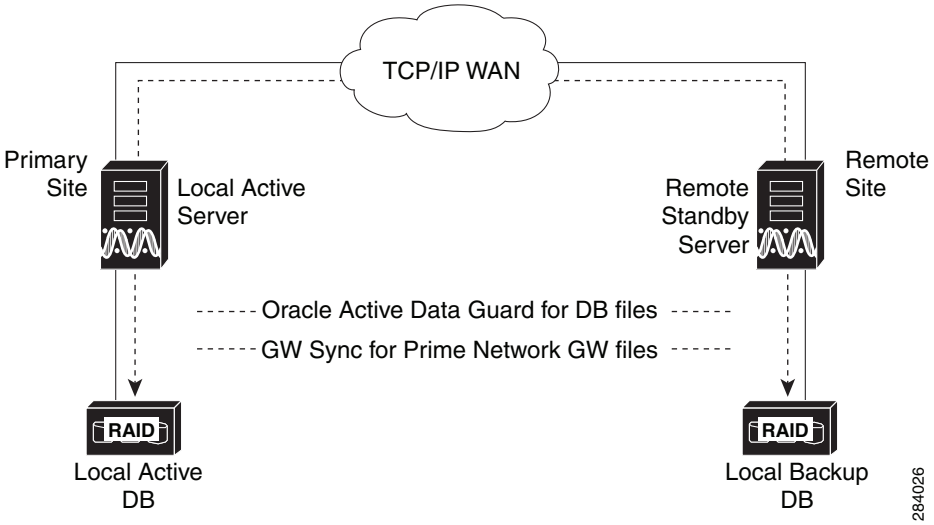

Note Geographical redundancy does not allow the Prime Network service (ana) to be brought online on the local side while the Oracle service is online on the remote side (or vice versa).

Oracle ADG Replication Process and Configuration Files

When the ADG solution is installed, a standby database is created at the remote site. The remote standby database is an active (read-only) Oracle instance. The local active database, which operates in archive log mode, sends copies of the redo logs to the standby 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 18-4 illustrates how data is replicated between the local active database and the remote standby database.

Figure 18-4 ADG Database Replication Process (ADG Geographical Redundancy)

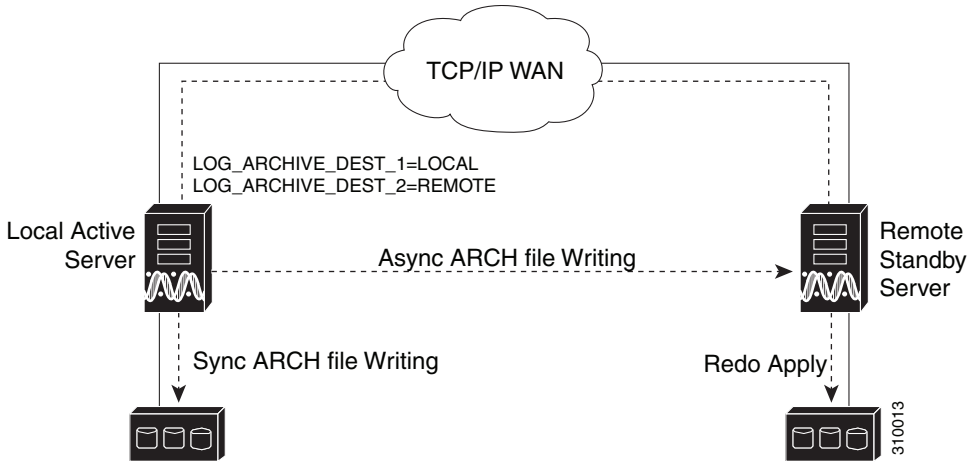

Note The databases must have identical disk capacities and mount points.

The following tables provide example parameters for the ADG int.ora and tnsnames.ora configuration files. These files reside on both the local and remote servers. In these examples, the local active database is named anadb and the remote standby database is named anadb_sb.

Table 18-2 Parameters for Oracle ADG int.ora Configuration Files

Parameter	Definition for int.ora at Local Active Site	Definition for int.ora Remote Standby Site
db_unique_name	anadb	anadb_sb
log_archive_dest_1	LOCATION=*log-archive-dest1-full-pathname*	
log_archive_dest_2	Service=anadb_sbASYNCLGWRVALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE) db_unique_name=anadb_sb	Service=anadb_sbASYNCLGWRVALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE) db_unique_name=anadb
log_archive_dest_state_1	enable	
log_archive_dest_state_2	enable	
standby_file_management	AUTO	

Table 18-3 shows example configuration parameters for the tnsnames.ora file at the local and remote sites. These files must be identical at both sites.

Table 18-3 Parameters for Oracle ADG tnsnames.ora Configuration Files

Parameter	Definition
ANADB=	(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = *ip-address*)(PORT = 1521)) (CONNECT_DATA = (SERVER = DEDICATED) (SERVICE_NAME = anadb)))
ANADB_SB=	(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = *ip-address*(PORT = 1521)) (CONNECT_DATA = (SERVER = DEDICATED) (SERVICE_NAME = anadb)))

To troubleshoot problems with the replication process, see the *Cisco Prime Network 3.8 Installation Guide*.

GWSync Replication Process

The GWSync utility is based on RHEL rsync. GWSync replicates the local primary server home directory (and any file system data that is required for disaster recovery) on the remote secondary server. Cron jobs trigger synchronization at both the primmer and secondary 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.

Figure 18-5 How GWSync Replication Process is Monitored (ADG Geographical Redundancy)

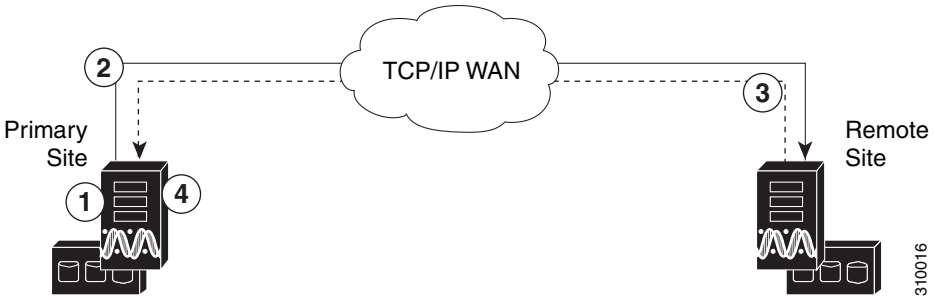

1	Local primary site generates local_timestamp file.	**3**	Primary site pulls remote site's timestamp file as remote_timestamp.
2	Remote secondary site pulls *NETWORKHOME* directory from local primary site (including remote site's local_timestamp file).	**4**	Primary site compares local_timestamp and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

To troubleshoot problems with the replication process, see the *Cisco Prime Network 3.8 Installation Guide*.

Embedded Database

Backups must be enabled for the embedded database, as described in Enabling Backups (Embedded Database), page 11-11.

IP Addresses and the network-conf Script

If you are using the **network-conf** script, when you are prompted for the IP address of units, use the floating IP address of the gateway.

File System Type

The ADG geographical redundancy solution requires the ext3 file system.

Security

To secure the channel used for data replication, an SSH key exchange is performed during the Prime Network installation.

LDAP External Authentication

If you use LDAP authentication in a geographical redundancy configuration, the gateway servers must be configured to communicate with two different LDAP servers, one at the local site and one at the remote site. For this reason the switchover and failover utilities will prompt you for the relevant LDAP parameters. The LDAP parameters are set once using Prime Network Administration.

If for some reason the necessary IP addresses are not updated after a switchover or failover, you can set them manually (which includes setting the necessary LDAP parameters). See Changing the Gateway IP Address on a Gateway and All Units (changeSite.pl), page 18-22.

For more information on using LDAP for user authentication, see Using an External LDAP Server for Password Authentication, page 7-7.

Licenses

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. See Licensing and Gateway Server High Availability, page 5-2.

Switchover/Failover Scenarios (ADG Geographical Redundancy)

These topics provide overviews of the switchover, failover, and fallback scenarios for ADG geographical redundancy configurations. The utilities used for these operations are stored in /var/adm/cisco/prime-network/scripts/ha/util.

Switchover and Fallback

A *switchover* is a planned, scheduled move from the primary active site to the secondary standby site when both sites are up. It is performed from the primary site using the **primeha -switch** command. (If local redundancy is configured at the primary site, it is performed from the node that contains the primary database.) The switchover reverses the replication direction for ADG and GWSync. If units are configured, the switchover script reconfigures the units to use the new active gateway and database.

A *fallback* is the process of reverting back to the original configuration. A fallback is also performed using the **primeha -switch** command, which causes the replication processes to revert back to their original direction.

These operations can only be performed from the primary site. For information on using the **primeha -switch** command, see Performing a Schedule Site Move (primeha - switch), page 18-15.

Failover and Fallback

A *failover* is normally the result of a serious failure which renders the primary site unavailable. In the case of such a failure, you must manually trigger a failover using the **primeha -fail** command, which disconnects the two sites, stops the replication process, and starts the standby server so it becomes a standalone node (that is, without geographical redundancy). These operations are performed from

standby site. (If local redundancy is configured at the standby site, it is performed from the node that contains the standby database.) For information on how to use **primeha**, see Managing ADG Geographical Redundancy, page 18-14.

Whether any data is lost depends on whether one of the sites is down when the failover occurs, because the failover event interrupts the replication process. If both sites are up, an orderly migration of data can be performed. Because replication channels are severed during the failover, you must reestablish all replication using the **setup_Prime_DR.pl** script (as described in the *Cisco Prime Network 3.8 Installation Guide*).

After all failures have been addressed and repaired, and replication is reinitiated, use the **primeha -switch** command to perform a fallback to the original setup.

These operations are performed from standby site. (If local redundancy is configured at the standby site, it is performed from the node that contains the standby database.) For information on how to use **primeha**, see Managing ADG Geographical Redundancy, page 18-14.

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

Recovering from a Disaster

For help recovering from a catastrophic failure, contact your Cisco account representative.

Managing ADG Geographical Redundancy

These topics provide information pertaining to ongoing management of an ADG geographical redundancy configuration.

Monitoring System Events

Prime Network generates the following System events for geographical redundancy monitoring:

- Informational event to indicate that both ADG and GWSync monitoring is active. This is done on an hourly basis based on cron jobs.
- Critical events when the following occur:
  - An GWSync has not occurred in the last 10 minutes.
  - The standby database is down.
  - The standby database is up but has been out of sync for 30 minutes.

To troubleshoot problems with the replication process, see the *Cisco Prime Network 3.8 Installation Guide*.

Monitoring Log Messages

The log files for data replication are described in the following table. To troubleshoot problems with the replication process, see the *Cisco Prime Network 3.8 Installation Guide*.

Log File	Description
NETWORKHOME/.replication *NETWORKHOME*/.replication_remote	Contains the local and remote timestamps used by GWSync.
NETWORKHOME/.replication_log	This log is only populated if the GWSync local and remote timestamps are more than 10 minutes apart (and a System event is generated), as in the following example: `Replication failed since: date`

NETWORKHOME/oracle_monitoring.log	Information on the Redo-apply log from the standby server.
	`+ Testing the replication state on the remote database`
	
	`- Redo transport lag:`
	`NAME VALUE TIME_COMPLETED`
	`---`
	`transport lag +00 00:00:00 04/14/2011 10:30:34`
	
	`- Redo apply lag:`
	`NAME VALUE TIME_COMPLETED`
	`---`
	`apply lag +00 00:00:00 04/14/2011 10:30:35`
	
	`- Active apply rate:`
	`ITEM UNITS SO FAR`
	`---`
	`Active Apply Rate KB/sec 286`
	
	`- Data base role: PHYSICAL STANDBY`

Checking Overall Status Using primeha

The **primeha** command is a central tool for checking the status of the high availability nodes, performing switchovers and failovers, and stopping and resuming data replication. The following example shows a configuration for a network that has both local and geographical redundancy.

- The first portion of the output shows the status of the geographical redundancy configuration. The server csi-exy.cisco.com is the remote standby gateway and database server. The server csi-w47.cisco.com is the other node in the local redundancy cluster and he is not running any service.
- The second portion of the output (that begins with Cluster Status) shows the status of the local redundancy configuration. (This is displayed because this setup also contains a local redundancy configuration.)

```
# perl primeha -status

+ Installing perl for HA
 - Installing ActivePerl-5.10.1.1007-x86_64-linux-glibc-2.3.3-291969
 - Extracting additional modules

 HOST                ANA SERVICE                                ORACLE SERVICE

 csi-bvc.cisco.com   Active Prime Network                     Active oracle    local
 csi-exy.cisco.com   Standby Prime Network                    Standby oracle
 csi-w47.cisco.com   Prime Network not running on this node    oracle not running on this
node

Cluster Status for ana_cluster @ Mon Aug  1 12:34:40 2011
Member Status: Quorate

 Member Name                     ID        Status
 ------ ----                     ----      ------
 csi-bvc.cisco.com               1         Online, Local, rgmanager
 csi-w47.cisco.com               2         Online, rgmanager

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

Performing a Schedule Site Move (primeha - switch)

Use the **primeha -switch** command to perform a scheduled move from a local primary site to a remote secondary site, when both sites are active. This is called a switchover. This is used for planned switches initiated by administrators.

The **primeha -switch** command will use the inputs you provided when you installed the gateway server high availability solution but will also give you an opportunity to modify those settings before performing the switchover. The switchover process consists of the following:

- Switch the roles between the primary and secondary sites.
- Switch the data replication sides (ADG and GWSync). In other words, the new primary site will be replicated to the new secondary site.

You can also use the switchover command to fallback to the primary site when a failed server is brought back online. The switchover will again reverse the replication directions. After performing a manual switchover, move any AVMs from unreachable units at the primary site to reachable units at the remote site.

Note This script must be run from the server with the *primary active* database.

To perform a switchover:

Step 1 Log into the server that contains the primary active database. (You can validate this by running **primeha -status**.)

Step 2 Move to the proper directory and start the script. The script will use the inputs you provided when you installed the gateway server high availability solution but will also give you an opportunity to modify those settings before performing the switchover.

Keep the following notes in mind:

- If the setup includes a dual-node cluster, when you are prompted for the gateway and database IP addresses, use the floating IP addresses for the Prime Network and Oracle services.
- You are only prompted for the "other cluster node" if the utility is invoked from a server that is part of a local redundancy setup. You should enter the IP address of the other cluster node—that is, the node the script is *not* being run from.

```
# cd /var/adm/cisco/prime-network/scripts/ha/util
# perl primeha -switch
+ Switching over to remote node
+ These are the parameters for the switchover process
  you will switch over to :
        gateway  : 10.56.56.57
        database : 10.56.56.57
        other cluster node : 10.56.56.67
        Prime Network user : network38
        Prime Network user home : /export/home/network38
        oracle user : oracle
        oracle user home : /opt/ora/oracle
```

Step 3 Approve or edit your switchover choices at the following prompt:

```
Do you approve? (yes/no)
```

- IIf you say **yes** and the system is using external authentication (LDAP), provide the necessary information at the following prompt (see Table 7-3 on page 7-11):

  ```
  Does this setup have an LDAP configured?
  ```

  Otherwise, proceed to Step 4.
- If you say **no**, you are prompted for the following information:

Field	Description
IP address of the remote gateway server	IP address of the standby gateway. If the remote node is a member of a dual-node cluster, use the floating IP address.
Root password for the node that has the gateway mounted	For the remote gateway server, the root password for the operating system (required for SSH).
IP address of the remote database	IP address of the standby database. If the remote node is a member of a dual-node cluster, use the floating IP address.
Root password for the node that has the database mounted	For the remote database, the root password for the operating system (required for SSH).

Field	Description
IP address of “other cluster node”	(If the local node is a member of a dual-node cluster) The IP address of the other node in the cluster.
Password for the other cluster node	(If the local node is a member of a dual-node cluster) For the other node in the cluster, the root password for the operating system (required for SSH).
Name for the OS user of the database	Name of database OS user.
Home directory of the user	Home directory for database OS user.
Name for the OS user for Prime Network	Name of Prime Network OS user.
Home directory of the user	Home directory for Prime Network OS user.
Whether the setup has LDAP configured	If system users LDAP (external authentication) for user authentication (see Table 7-3 on page 7-11).

Step 4 Confirm that you want to continue with the switchover. Prime Network proceeds and displays text similar to the following.

```
- Checking if Prime Network is mounted on local node                    [MOUNTED]
- Verifying local oracle status
- Verifying remote oracle status
- Stopping Prime Network on local side..                                 [OK]
- Removing replication monitoring cron                                  [OK]
- Changing local Prime Network flag to remote                           [OK]
- Copying scripts to sub
- Switching local server to remote
- Changing local oracle flag to remote
- Copying scripts to remote database
- Running switchover script on remote database
- Copying scripts to remote gateway
- Running switchover script on remote gateway
- Switching local server to recover mode
- Set db to read only mode
```

Step 5 If required, manually move the AVMs from the unreachable units at the primary site to the reachable units at the remote site. See Moving and Deleting AVMs, page 4-13. (This is not required if the local units were not affected by a failure; the script will reconfigure the units to use the relevant gateway and database.)

Step 6 Verify that the new gateway IP address and database IP addresses are correct. If needed, switch the IP address manually using one of the following procedures:

- Changing the Gateway IP Address on a Gateway and All Units (changeSite.pl), page 18-22
- Changing the Gateway IP Address on a Single Unit (switchUnit.pl), page 18-24

Using Failover for Disaster Recovery (primeha -fail)

Caution Failover is time-consuming and requires the system to be shut down. It should only be used when the primary site fails. Do not execute it until all other options for restoring the primary site are explored.

Note A *manual failover* should only be performed when the primary site has failed.

Use the **primeha -fail** command to perform a site switch for disaster recovery. A site switch is a manual move from the local primary site to the remote secondary site. The script will use the inputs you provided when you installed the gateway server high availability solution but will also give you an opportunity to modify those settings before performing the failover. When you invoke **primeha -fail**, the command does the following:

- Disconnects the primary site from the secondary site.
- Stops the GWSync and ADG replication processes.
- Start the standby server as standalone node without geographical redundancy.

After performing a manual failover, move any AVMs from unreachable units at the primary site to reachable units at the remote site.

Note The failover must be run from the node that contains the *standby* database. If the system is using external authentication (LDAP), you will have to provide the LDAP URL, distinguished name prefix and suffix, and the protocol (see Table 7-3 on page 7-11).

To perform a failover:

Step 1 As root, log into the active node that contains the standby database. (You can validate this by running **primeha -status**.)

Move to the proper directory and start the script. The script will use the inputs you provided when you installed the gateway server high availability solution but will also give you an opportunity to modify those settings before performing the failover.

Keep the following notes in mind:

- If the setup includes a dual-node cluster, when you are prompted for the gateway and database IP addresses, use the floating IP addresses for the Prime Network and Oracle services.
- You are only prompted for the “other cluster node” if the utility is invoked from a server that is part of a local redundancy setup. You should enter the IP address of the other cluster node—that is, the node the script is *not* being run from.

```
# cd /var/adm/cisco/prime-network/scripts/ha/util
# perl primeha -fail
+ Failing over to remote node
+ These are the parameters for the fail over process
  you will fail over to :
        gateway  : 10.56.56.74
        database : 10.56.56.41
  from :
        gateway  : 10.56.56.57
        database : 10.56.56.57
        other cluster node : 10.56.56.67
        Prime Network user : network38
```

```
Prime Network user home : /export/home/network38
oracle user : oracle
oracle user home : /opt/ora/oracle
```

Step 2 Approve or edit your switchover choices at the following prompt:

```
Do you approve? (yes/no)
```

- If you say **yes** and the system is using external authentication (LDAP), provide the necessary information at the following prompt (see Table 7-3 on page 7-11):

  ```
  Does this setup have an LDAP configured?
  ```

  If you say **yes** and the system is *not* using external authentication, proceed to Step 3.

- If you say **no**, you are prompted for the following information:

Field	Description
IP address of the remote gateway server	IP address of the standby gateway. If the remote node is a member of a dual-node cluster, use the floating IP address.
Root password for the node that has the gateway mounted	For the remote gateway server, the root password for the operating system (required for SSH).
IP address of the remote database	IP address of the standby database. If the remote node is a member of a dual-node cluster, use the floating IP address.
Root password for the node that has the database mounted	For the remote database, the root password for the operating system (required for SSH).
IP address of “other cluster node”	(If the local node is a member of a dual-node cluster) The IP address of the other node in the cluster.
Password for the other cluster node	(If the local node is a member of a dual-node cluster) For the other node in the cluster, the root password for the operating system (required for SSH).
Name for the OS user of the database	Name of database OS user.
Home directory of the user	Home directory for database OS user.
Name for the OS user for Prime Network	Name of Prime Network OS user.
Home directory of the user	Home directory for Prime Network OS user.
Whether the setup has LDAP configured (yes/no)	If system users LDAP (external authentication) for user authentication (see Table 7-3 on page 7-11).

Step 3 Confirm that you want to continue with the failover. Prime Network proceeds and displays text similar to the following.

```
- Checking if Prime Network is mounted on local node                    [MOUNTED]
 - Verifying local oracle status
 - Copying scripts to remote gateway
 - Running failover script on remote gateway
 - Copying scripts to remote database
 - Running failover script on remote database
 - Switching local db to active mode
 - Changing remote oracle flag to local
 - Starting replication monitoring cron                                 [OK]
 - Changing remote Prime Network flag to local                          [OK]
```

```
- Copying scripts to sub
- Running script on cluster standby node
```

Step 4 Move any AVMs from unreachable units at the primary site to reachable units at the remote site. See Moving and Deleting AVMs, page 4-13.

Step 5 Verify that the new gateway IP address and database IP addresses are correct. If needed, switch the IP address manually using one of the following procedures:

- Changing the Gateway IP Address on a Gateway and All Units (changeSite.pl), page 18-22
- Changing the Gateway IP Address on a Single Unit (switchUnit.pl), page 18-24

Stopping Data Replication (primeha -stop)

Use the stop replication command **primeha -stop** when you need to perform scheduled work on a server in the remote site. It stops the replication process to the remote site and shuts down the remote database. To resume replication, see Resuming Data Replication (primeha -start), page 18-21.

Note This command *must* be run from the server that contains the standby database. (You can validate this by running **primeha -status**.)

To run the stop replication command:

```
# cd /var/adm/cisco/prime-network/scripts/ha/util
# perl primeha -stop
```

The following is an example of a stop replication session. In this example:

- The remote standby gateway and database IP address is 10.56.56.57. The user wants to stop the replication of data from the local site to this remote site.
- The local gateway IP address is 10.56.57.74 and the local database IP address is 10.56.56.41.

This utility must be run from the server with the remote standby database (in this example, 10.56.56.57). This will stop replicating data and will shut down the remote database.

Keep these notes in mind when you are prompted for the following information:

- Remote server's gateway IP address—Enter the IP address for the primary gateway. If the primary site has a local redundancy setup, enter the floating IP address for the Prime Network service.
- Remote data base IP address—Enter the IP address for the primary database. If the primary site has a local redundancy setup, enter the floating IP address of the Oracle service.
- Cluster sub server's IP address—(Is displayed only if the standby database is part of a cluster) Enter the physical IP address of the other cluster node—that is, the node the script is *not* being run from.

```
[root@10.56.56.57]# perl primeha -stop

+ Installing perl for HA
- Installing ActivePerl-5.10.1.1007-x86_64-linux-glibc-2.3.3-291969
- Extracting additional modules

+ Stopping replication to remote node
- Enter the remote server's gateway IP address:
10.56.56.74
- Enter the root password for the node that has the gateway mounted
- Enter the remote data base IP address:
10.56.56.41
```

```
- Enter the root password for the node that has the data base mounted
- Enter a name for the OS user of the database [oracle]
- Enter the home directory of the user (oracle) [/opt/ora/oracle]
- Enter a name for the OS user for Prime Network [network38]
- Enter the home directory of the user (Prime Network) [/export/home/network38]
- Checking if Prime Network is mounted on local node          [MOUNTED]
- Removing local node Prime Network flag
- Stopping local db replication
- Removing local node data base flag
- Stopping replication on remote gateway
- Copying scripts to remote database
- Running stop replication script on remote database

+ Removing perl for HA
```

Resuming Data Replication (primeha -start)

Note This command can only be used if (1) the remote database was stopped using **primeha -stop**, and (2) the remote database has *not* been down for more than seven days. If the remote database *has* been down for more than seven days, you must recreate the remote database by using the **setup_Prime_DR.pl** script (see the *Cisco Prime Network 3.8 Installation Guide* for information on how to use that script.)

Use the resume replication utility **primeha -start** to start the database at the remote site (in open, read-only mode) and restart the replication process. This command should only be used after stopping replication in order to perform scheduled work on the remote site.

Note This command must be run from the server that contains the remote standby database. (You can validate this by running **primeha -status**.)

Keep these notes in mind when you are prompted for the following information:

- Remote server's gateway IP address—Enter the IP address for the primary gateway. If the primary site has a local redundancy setup, enter the floating IP address for the Prime Network service.
- Remote data base IP address—Enter the IP address for the primary database. If the primary site has a local redundancy setup, enter the floating IP address of the Oracle service.
- Cluster sub server's IP address—(Is displayed only if the standby database is part of a cluster) Enter the physical IP address of the other cluster node—that is, the node the script is *not* being run from.

To run the resume replication command:

```
# cd /var/adm/cisco/prime-network/scripts/ha/util
# perl primeha -start
```

The following is an example of a **primeha -start** session. It uses the same parameters as the stop replication example:

- The remote standby gateway and database IP address is 10.56.56.57. The user wants to restart the replication of data from the local site to this remote site.
- The primary gateway IP address is 10.56.57.75 and the primary database IP address is 10.56.56.41.

The user wants to restart the remote database in read-only mode (in other words, make it the standby database), and resume replicating data. This utility is run from the node with the remote standby database (10.56.56.57).

```
[root@10.56.56.57]# perl primehah -start

+ Installing perl for HA
- Installing ActivePerl-5.10.1.1007-x86_64-linux-glibc-2.3.3-291969
- Extracting additional modules

+ Resuming replication to remote node
- Enter the remote server's gateway IP address:
10.56.56.74
- Enter the root password for the node that has the gateway mounted
- Enter the remote data base IP address:
10.56.56.41
- Enter the root password for the node that has the data base mounted
- Enter a name for the OS user of the database [oracle]
- Enter the home directory of the user (oracle) [/opt/ora/oracle]
- Enter a name for the OS user for Prime Network [network38]
- Enter the home directory of the user (Prime Network) [/export/home/network38]
- Resuming local db replication
- Adding local node data base flag
- Checking if Prime Network is mounted on local node          [MOUNTED]
- Adding local node Prime Network flag
- Resuming replication on remote gateway
- Copying scripts to remote database
- Running resume replication script on remote database

+ Removing perl for HA
```

Changing IP Addresses (ADG Geographical Redundancy)

If all IP addresses are not automatically changed after a failover or switchover, use the following procedures, as appropriate.

- Changing the Gateway IP Address on a Gateway and All Units (changeSite.pl), page 18-22
- Changing the Gateway IP Address on a Single Unit (switchUnit.pl), page 18-24

Changing the Gateway IP Address on a Gateway and All Units (changeSite.pl)

If the gateway IP address is not updated on any of the units (or on the gateway) during a site-to-site failover or switchover, use the **changeSite.pl** utility to do so manually. This procedure will change the address on the gateway and all reachable units.

Note If a dual-node cluster is part of a local redundancy setup, use the logical IP addresses.

The following table describes the options to the **changeSite.pl** utility. If you are using an external LDAP server for user authentication, you must also set the necessary LDAP parameters, as described below. For more details on these parameters, see Configuring Prime Network to Communicate with the External LDAP Server, page 7-10.

Option	Description
-force	Allow manual change to registry settings. (Because this script runs as part of a failover or switchover, the **-force** option is required when running the script from the command line.)
-newgwip *new-gateway-ip*	IP address of the gateway that is running after the failover or switchover.
-newdbip *new-database-ip*	IP address of the database that running after the failover or switchover.
-oldgwip *old-gateway-ip*	IP address of the gateway that was running prior to the failover or switchover.
-oldbip *old-database-ip*	IP address of the database that was running prior to the failover or switchover.
[**-newldapurl** *new-ldap-url* **-oldldapurl** *old-ldap-url*]	(LDAP only) URL for the LDAP server that will be used by the running gateway (*new-ldap-url*), and the URL that was used by the gateway that was running prior to the failover or switchover (*old-ldap-url*). Use the following format: **ldap://***host.company***.com:***port*
[**-newldapprefix** *new-ldap-prefix* **-oldldapprefix** *old-ldap-prefix*]	(LDAP only) First part of the LDAP DN (which is used to uniquely identify users) for the new and old LDAP server. Both *new-ldap-prefix* and *old-ldap-prefi*x should be entered exactly as shown below: **CN** (The actual format is **CN=***Value*, which specifies the common name for specific users. =*Value* will be automatically populated with Prime Network usernames.)
[**-newldapsuffix** *new-ldap-suffix* **-oldldapsuffix** *old-ldap-suffix*]	(LDAP only) Second part of the LDAP distinguished name, which specifies the location in the directory for both the new and old LDAP servers. Both *new-ldap-suffix* and *old-ldap-suffi*x should use the following format **,CN=Users,DC=***LDAP_server***,DC=***company***,DC=com**
[**-newldapisssl** *new-ldap-is-ssl* **-oldldapisssl** *old-ldap-is-ssl*]	(LDAP only) Encryption protocol to be used for communication between the running Prime Network gateway server and the new LDAP server (*new-ldap-is-ssl*), and the protocol that was used between the old gateway and LDAP servers (*old-ldap-is-ssl*).

Step 1 If you will reset LDAP information, reconfigure them first from the Prime Network Administration GUI client. See Configuring Prime Network to Communicate with the External LDAP Server, page 7-10.

Step 2 Log into the primary gateway server as *network-user*.

Step 3 Change to the correct directory:

```
# cd $ANAHOME/Main/ha
```

Step 4 Run the following command:

```
# perl changeSite.pl -force -newgwip new-gw-ip -newdbip new-db-ip
 -oldgwip old-gw-ip -olddbip old-db-ip
[-newldapurl new-ldap-url -oldldapurl old-ldap-url]
[-newldapprefix new-ldap-prefix -oldldapprefix old-ldap-prefix]
[-newldapsuffix new-ldap-suffix -oldldapsuffix old-ldap-suffix]
[-newldapisssl new-ldap-is-ssl -oldldapisssl old-ldap-is-ssl]
```

The following is an example of a **changeSite.pl** session. In this example the following is being changed:

- The original gateway and database IP address was 10.56.56.57.
- The site was switched over to the standby gateway (10.56.56.74) and database (10.56.56.41).

For some reason, the IP addresses were not correctly changed to reflect the new addresses. The utility forces the IP addresses to be changed to 10.56.56.74 for the gateway and 10.56.56.41 for the database. In this example the system is not using LDAP, so those parameters are not included.

```
csi-exy% cd $ANAHOME/Main/ha
csi-exy% perl changeSite.pl -force -newgwip 10.56.56.74 -newdbip 10.56.56.41 -oldgwip
10.56.56.57 -olddbip 10.56.56.57
Thu Apr 14 16:08:22 2011 --[INFO]: '-Forced change of gw address from 10.56.56.57 to
10.56.56.74.... '
Thu Apr 14 16:08:22 2011 --[INFO]: '--changing uplinks for gw AVM0'
Thu Apr 14 16:08:22 2011 --[INFO]: '--changing uplinks for unit AVM0s'
Thu Apr 14 16:08:22 2011 --[INFO]: '--changing gw ip and haservice for unit AVM99s'
Thu Apr 14 16:08:22 2011 --[INFO]: '--changing registry on units'
Thu Apr 14 16:08:37 2011 --[INFO]: '--changing localhost entry for gw AVM99'
Thu Apr 14 16:08:37 2011 --[INFO]: '-Forced change of db server address from 10.56.56.57
to 10.56.56.41.... '
Thu Apr 14 16:08:37 2011 --[INFO]: '--changing db server ip for gw AVM66'
Thu Apr 14 16:08:38 2011 --[INFO]: '--changing db server ip for gw persistency.xml'
Thu Apr 14 16:08:38 2011 --[INFO]: '--changing db server ip for template persistency.xml'
Thu Apr 14 16:08:38 2011 --[INFO]: '--changing db server ip for unit persistency.xml'
Thu Apr 14 16:08:38 2011 --[INFO]: '-Forced change of NCCM address from 10.56.56.57 to
10.56.56.41.... '
new IP address is: 10.56.56.41
jdbc.properties file has been updated to change to new IP address
Thu Apr 14 16:08:39 2011 --[INFO]: '->Done'
```

Changing the Gateway IP Address on a Single Unit (switchUnit.pl)

If any of the units do not reflect the updated gateway and database IP address after a site-to-site failover or switchover, use the **switchUnit.pl** utility to do so manually. This procedure will change the address only on the unit from which it is run.

Note If a dual-node cluster is part of a local redundancy setup, use the logical IP addresses.

For any unit that does not reflect the updated gateway and database IP addresses:

Step 1 Log into the unit as *network-user*.

Step 2 Change to the correct directory:

```
# cd $ANAHOME/Main/ha
```

Step 3 Run the following command:

```
# perl switchUnit.pl new-gw-ip old-gw-ip new-db-ip old-db-ip
```