



Performing Switchovers and Failovers

These topics describe how to use the high availability scripts to perform a switchover between two active sites, perform a failover when a primary site fails, restore the configuration of a failed site, and other high availability operations.

- [Performing a Scheduled Site Move, page 6-1](#)
- [Failing Over to the Standby Site for Disaster Recovery, page 6-4](#)
- [Restoring the Failed Site \(Hot Backup\), page 6-6](#)
- [Stopping and Restarting Data Replication, page 6-11](#)
- [Changing IP Addresses after a Failover or Switchover, page 6-14](#)

Performing a Scheduled Site Move

Use the **primeha -switch** command to perform a scheduled move from a primary site to a remote site, when both sites are available. 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 standby sites.
- Switch the data replication sides (ADG and GWSync). In other words, the new primary site will be replicated to the new standby site.
- For Operations Reports, start AVM 45 on the standby 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

primeha -switch command 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.

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

Make sure:

- If switching to a dual-node cluster, when you are prompted for the gateway and database IP addresses, use the floating IP addresses for the Prime Network, Oracle, and Infobright (if installed) services.
- You are only prompted for the “remote 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 remote node—that is, the node the script is *not* being run from.

```
+ Switching over to remote node
+ These are the parameters for the switchover process
you will switch over to :
  gateway   : 1.1.1.1
  database  : 1.1.1.1
  remote node : 1.1.1.2
  Prime Network user : pn41
  Prime Network user home : /export/home/pn41
  oracle user : oracle
  oracle user home : /opt/ora/oracle
```

Step 3 Approve or edit the default choices that appear based on the inputs you provided when you installed the gateway server high availability:

Do you approve? (yes/no)

- If you enter **yes** and the system is using external authentication (LDAP), provide the necessary information at the following prompt. For details on these LDAP prompt, see [Configuring Prime Network to Communicate with the External LDAP Server in the *Cisco Prime Network 4.3 Administrator Guide*](#).

Does this setup have an LDAP configured? (yes/no)

Otherwise, proceed to [Step 4](#).

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

Field	Description
IP address of the remote gateway server on which Prime Network will run.	IP address of the standby gateway.
Root password for the remote node on which Prime Network will run.	For the remote site gateway server, the root password for the operating system (required for SSH).
IP address of the remote server\service on which Oracle database will run	IP address of the standby database.
Root password of the remote server on which Oracle database will run	For the remote database, the root password for the operating system (required for SSH).

Field	Description
Name for the OS user of the database	Example: Oracle
Home directory of the user	Example: /opt/ora/oracle
Name for the OS user for Prime Network	Example: <i>pn41</i>
Home directory of the user	Example: /export/home/ <i>pn41</i>
Whether the setup has LDAP configured	If system users LDAP (external authentication) for user authentication (see the Cisco Prime Network 4.3 Administrator Guide).

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
- Changing local Prime Network flag to remote                 [OK]
- Stopping Prime Network on local side..                      [OK]
- Switching local server to remote
- Copying scripts to remote database
- Running pre-switchover script on remote database
- 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. For moving and deleting AVMs information see the [Cisco Prime Network 4.3 Administrator Guide](#). (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 6-14
- [Changing the Gateway IP Address on a Single Unit \(switchUnit.pl\)](#), page 6-16

Step 7 To verify the setup, perform all of the tests (not including the step for creating database links) that are described in [Verifying the Geographical Redundancy Setup](#), page 4-14.



Note

To make cross-launch work on the upgrade setup, de-register and register the Prime Network, after the Prime Central and Prime Network switchover.

Performing Switchover on Systems with Prime Network Integration Layer Installed on Prime Network

Use the below procedure to perform switch over from local node (also local cluster node) to remote (DR) server (S2) with Prime Network Integration Layer installed.

To perform a switchover on systems that has PN-IL installed on top of Prime Network.

-
- Step 1** As root user, log in to Prime Network primary server and perform switch over to DR node, using the procedure in [Performing a Scheduled Site Move, page 6-1](#).
- Step 2** After switchover, login to the DR node as the Prime Network user.
- ```
su - pnuser
```
- Step 3** Enable PN-IL Health monitor. [Health monitor will bring up PN-IL service if it is down.]
- ```
$PRIMEHOME/local/scripts/il-watch-dog.sh enable
```
- Step 4** After the primary server is up, login as a root user to the Prime Network DR node and again perform the switch over to primary node, using the procedure in [Performing a Scheduled Site Move, page 6-1](#).
- Step 5** After switchover, log in to the DR node as the Prime Network user and disable the PN-IL health monitor
- ```
$PRIMEHOME/local/scripts/il-watch-dog.sh disableandstop
```
- 

## Failing Over to the Standby Site for Disaster Recovery



### Note

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

Use the **primeha -fail** command to perform a site failover for disaster recovery. A site failover is a manual move from the failed primary site to the standby site at remote location. 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 remote 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.

If you are using Operations Reports, the data from the past 1 hour and 20 minutes will be lost.



### 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 [Configuring Prime Network to Communicate with the External LDAP Server in the Cisco Prime Network 4.3 Administrator Guide](#)).

To perform a failover:

- 
- Step 1** As a root user, 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.

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

Make sure:

- 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 “remote 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 remote node—that is, the node the script is *not* being run from.

```
+ Failing over to remote node
+ These are the parameters for the fail over process
you will fail over to :
 gateway : 1.1.1.1
 database : 1.1.1.2
from :
 gateway : 1.1.1.1
 database : 1.1.1.1
 remote node : 1.1.1.3
 Prime Network user : pn41
 Prime Net work user home : /export/home/pn1
 oracle user : oracle
 oracle user home : /opt/ora/oracle
```

**Step 2** Approve or edit the default choices that appear based on the inputs you provided when you installed the gateway server high availability:

Do you approve? (yes/no)

- If you enter **yes** and the system is using external authentication (LDAP), provide the necessary information at the following prompt. For details on these LDAP prompt, see [Configuring Prime Network to Communicate with the External LDAP Server in the \*Cisco Prime Network 4.3 Administrator Guide\*](#):

Does this setup have an LDAP configured?

If you enter **yes** and the system is *not* using external authentication, proceed to [Step 3](#).

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

| Field                                                           | Description                                                                                                                                                    |
|-----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP address of the server to which the gateway should fail over  | IP address of the standby gateway. If the remote site is a member of a dual-node cluster, use the floating IP address (of the management port of the cluster). |
| IP address of the server to which the database should fail over | IP address of the standby database. If the remote site is a member of a dual-node cluster, use the floating IP address.                                        |
| IP address of the current gateway server                        |                                                                                                                                                                |
| 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 current database server                       |                                                                                                                                                                |
| 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                                                                                                                                                                                                                                          |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OS user of the database                        | Name of database OS user.                                                                                                                                                                                                                            |
| Home directory of the user (oracle)            | Example: /opt/ora/oracle                                                                                                                                                                                                                             |
| OS user for Prime Network                      | Example: pn41                                                                                                                                                                                                                                        |
| Home directory of the user (pn41)              | Example: /export/home/pn41                                                                                                                                                                                                                           |
| Whether the setup has LDAP configured (yes/no) | If system users LDAP (external authentication) for user authentication (For details on LDAP, see <a href="#">Configuring Prime Network to Communicate with the External LDAP Server in the <i>Cisco Prime Network 4.3 Administrator Guide</i></a> ). |

**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. For moving and deleting AVMs information see the [Cisco Prime Network 4.3 Administrator Guide](#).

**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 6-14
- [Changing the Gateway IP Address on a Single Unit \(switchUnit.pl\)](#), page 6-16

**Note**

To restore the configuration after a disaster, see [Restoring the Failed Site \(Hot Backup\)](#), page 6-6.

**Step 6** To verify the setup, perform all of the tests (not including the step for creating database links) that are described in [Verifying the Geographical Redundancy Setup](#), page 4-14.

**Step 7** Log into Prime Network using the new IP address.

## Restoring the Failed Site (Hot Backup)

**Note**

In this section, *failed site* refers to the non-active site. This could be:

- A cluster (Server P1 and Server P2 in a local or geographic redundancy setup), or

- A single node (Server P1 in a geographical redundancy (only) setup).

After the servers are up and running on the failed site, use the procedures in this section to restore the redundancy configuration on the failed site. For information about how to failover to a standby site after a disaster, see [Failing Over to the Standby Site for Disaster Recovery](#), page 6-4.

While restoring the redundancy configuration on the failed site, you do not have to take down the active site. Both sites are up while the failed site resumes, and you can switch back to the active site, without any down time. Restoring the redundancy configuration on the failed site depends on whether the servers on the failed site were down due to a catastrophic or non-catastrophic failure.

- A catastrophic failure is when the servers becomes unoperative and unreachable, for example, unexpected breakdown on the server side due to disk crash, or sudden power surge. In this case, the existing setup or configuration is completely lost.
- A non-catastrophic failure is when the server may be unoperative for some time but the setup is still intact, for example, in case of a reboot.

The `resumeFromFailOver.pl` script is used for restoring the redundancy configuration and has the following format:

```
perl resumeFromFailOver.pl -setup_replication [-daemonize] | -reconfigure_setup | [-autoconf dir]
-reinstall_setup
```

Table 6-1 describes the arguments and option and also indicates the node from where the commands should be executed.

**Table 6-1 Options/Arguments for resumeFromFailOver.pl**

| Arguments/Options                                             | Description                                                                                                                                                                                                                                                                                                                                           | Failed Site <sup>1</sup> | Active Site |
|---------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|-------------|
|                                                               |                                                                                                                                                                                                                                                                                                                                                       |                          |             |
| <code>[-autoconf dir]</code><br><code>-reinstall_setup</code> | Reinstalls the gateway on the failed site <sup>1</sup> that can be either a dual node cluster or a single node server.<br><br>The <code>-autoconf dir</code> option runs the operation in non-interactive mode using input from the <code>rf_auto_install_RH.ini</code> file located in <code>dir</code> ( <code>dir</code> must be a full pathname). | x                        | -           |
| <code>-setup_replication</code><br><code>[-daemonize]</code>  | Restores the replication between the failed site and remote site. The <code>daemonize</code> option runs the replication process in background without any user interaction.                                                                                                                                                                          | -                        | x           |
| <code>-reconfigure_setup</code>                               | Reconfigures the failed <sup>2</sup> site <sup>1</sup> after a failover. Use this flag only when the setup still exists on the failed site.<br><br>If for some reason, the <code>--setup_replication</code> fails, then use this flag on the failed site first, and then run <code>--setup_replication</code> again (from the active site).           | x                        | -           |

1. The site which had failed. This could be a Cluster (Server P1 and Server P2 in a local or geographic redundancy setup) or only a single node (Server P1 in a geographical redundancy (only) setup)
2. Failed server still operates and Prime Network installation exists. Redundancy setup can be reconfigured without reinstalling.

Depending on the type of failure on the failed site, do one of the following to restore the redundancy configuration on the failed site:

- [Restoring Redundancy Configuration After a Catastrophic Failure, page 6-8](#)
- [Restoring Redundancy Configuration After a Non-Catastrophic Failure, page 6-10](#)

## Restoring Redundancy Configuration After a Catastrophic Failure

### Before You Begin:

Make sure all the installation and high availability requirements are met. See [Installation Requirements for Geographical Redundancy, page 4-4](#).



### Note

Ensure that you have already performed the failover procedure before proceeding to the restore redundancy configuration on the failed site. For information about how to failover to a standby site after a disaster, see [Failing Over to the Standby Site for Disaster Recovery, page 6-4](#).

To restore the redundancy configuration on the failed site after a catastrophic failure, do the following:

- Step 1** As a root user, log into the failed site and unzip the RH\_ha.zip located on the “Disk 1 New Install” DVD. Unzipping **RH\_ha.zip** creates the /tmp/RH\_ha directory. Also, unzip the **RH\_ha.zip** in the primary location.
- Step 2** From the /tmp/RH\_ha directory, run the **perl resumeFromFailOver.pl -reinstall\_setup** for reinstalling the failed site. For information about other options or arguments used with **resumeFromFailOver.pl** script, see [Table 6-1](#).
- ```
cd /tmp/RH_ha
perl resumeFromFailOver.pl -reinstall_setup
```
- Step 3** Enter y at the prompt to continue with the Prime Network installation.
- ```
Would you like to continue? (y/n)
```
- Step 4** Enter the server details as shown in [Table 6-2](#) or [Table 6-3](#), depending on whether it is local and geographical redundancy configuration or a geographical redundancy only.

**Table 6-2 Prompts that Appear While Restoring Local and Geographical Redundancy Configuration**

| Field                                               | Description                                                                                                                                                                                                                                       |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hostname of the primary node                        | Hostname of the active site, i.e, the site that is currently running both the cluster services (ana, oracle_db).                                                                                                                                  |
| Password for the primary node root user             | Password of the active site, i.e, the site that is currently running both the cluster services (ana, oracle_db).                                                                                                                                  |
| Whether NTP is configured on the 2 gateways (local) | 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 <a href="#">Cisco Prime Network 4.3 Installation Guide</a> . |
| Checking whether multipath is configured            | Check whether multipath is enabled.                                                                                                                                                                                                               |



**Table 6-2 Prompts that Appear While Restoring Local and Geographical Redundancy Configuration (continued)**

| Field                                                                 | Description                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Checking whether to run automated backup for embedded database backup | Indicates whether to run the embedded database automated backups, yes or no.                                                                                                                                                                                                                                        |
| Password for <i>pnuser</i>                                            | User-defined password for the <i>pnuser</i> .                                                                                                                                                                                                                                                                       |
| Location of the Prime Network build image                             | Enter full path to the installation image version that was first installed.<br><br>For example, if you had installed 3.8 first, then upgraded to 3.10, then upgraded to 4.1, and then upgraded to 4.2, 4.2.3, or 4.3 provide the path of the 3.8 image.                                                             |
| Directory for the oracle zip files                                    | The directory containing the embedded Oracle zip files.                                                                                                                                                                                                                                                             |
| Password for <i>NODE_ONE_HOSTNAME</i> root user                       | The root user password for the node running the installation. For local redundancy dual-node clusters, this node must be one of the cluster nodes.                                                                                                                                                                  |
| Hostname for the cluster second node                                  | Hostname of the second node in the cluster                                                                                                                                                                                                                                                                          |
| Password for <i>NODE_TWO_HOSTNAME</i> root user                       | Password of the second node in the cluster                                                                                                                                                                                                                                                                          |
| SMTP server                                                           | The local mail server                                                                                                                                                                                                                                                                                               |
| User email address                                                    | The email address to which error messages will be sent from the embedded database if problems occur.                                                                                                                                                                                                                |
| IP address for the ORACLE service                                     | Floating or Virtual IP of the Oracle service group.                                                                                                                                                                                                                                                                 |
| IP address for the Prime Network service                              | Floating or Virtual IP of the Prime Network service.                                                                                                                                                                                                                                                                |
| Valid multicast address for the cluster nodes                         | An available multicast address accessible and configured for both cluster nodes.                                                                                                                                                                                                                                    |
| Name for the Prime Network cluster                                    | User-defined cluster name. The cluster name cannot be more than 15 non-NUL (ASCII 0) characters. For local redundancy, the cluster name must be unique within the LAN.                                                                                                                                              |
| <i>NODE_ONE_HOSTNAME</i> fencing agent                                | The hostname of the fencing device configured for the node running the installation. This can be an IP address for some fencing devices.                                                                                                                                                                            |
| <i>NODE_TWO_HOSTNAME</i> fencing agent                                | The hostname of the fencing device configured for the second cluster node. For some fencing devices, this can be an IP address.                                                                                                                                                                                     |
| Password for the Prime Network cluster web interface admin user       | Indicates the port and the password for the cluster web interface. The <i>LUCI_PORT</i> must be available and should not be in the Prime Network debug range (60000 <= X < 61000) or in the Prime Network avm port range (2000 <= X < 3000) OR (8000 <= X < 9000). The password must contain at least 6 characters. |
| Port for the Prime Network cluster web interface                      |                                                                                                                                                                                                                                                                                                                     |

**Table 6-3 Prompts that Appear While Restoring Geographical Redundancy Configuration**

| Field                                                          | Description                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hostname of the primary node                                   | Hostname of the active server.                                                                                                                                                                                                                                  |
| Password for the primary node root user                        | The root user password for the active server.                                                                                                                                                                                                                   |
| Whether NTP is configured on the 2 gateways (local and remote) | 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 <a href="#">Cisco Prime Network 4.3 Installation Guide</a> .               |
| Password for <i>pnuser</i>                                     | User-defined password for the <i>pnuser</i> .                                                                                                                                                                                                                   |
| Location of the Prime Network build image                      | Enter full path to the installation image version that was first installed.<br><br>For example, if you had installed 3.8 first, then upgraded to 3.10, then upgraded to 4.1, and then upgraded to 4.2, 4.2.2, 4.2.3, or 4.3, provide the path of the 3.8 image. |
| Directory for the oracle zip files                             | The directory containing the embedded Oracle zip files.                                                                                                                                                                                                         |
| Password for <i>HOSTNAME</i> root user                         | The root user password for the host machine.                                                                                                                                                                                                                    |
| SMTP server                                                    | The local mail server                                                                                                                                                                                                                                           |
| User email address                                             | The email address to which error messages will be sent from the embedded database if problems occur.                                                                                                                                                            |

- Step 5** Log into the active site and run the **perl resumeFromFailOver.pl -setup\_replication** from the directory where **RH\_ha.zip** was extracted. This script will set up replication between the failed and the active site. For information on other options or arguments used with **resumeFromFailOver.pl** script, see [Table 6-1](#).
- Step 6** Confirm that you want to continue with setting up of replication process.
- Step 7** To verify the setup, perform all of the tests (not including the step for creating database links) that are described in [Verifying the Geographical Redundancy Setup, page 4-14](#).

## Restoring Redundancy Configuration After a Non-Catastrophic Failure

### Before You Begin:

Make sure all the installation and high availability requirements are met. See [Installation Requirements for Geographical Redundancy, page 4-4](#).



### Note

Ensure that you have already performed the failover procedure before proceeding to restore the redundancy configuration on the failed site. For information about how to failover to a standby site after a disaster, see [Failing Over to the Standby Site for Disaster Recovery, page 6-4](#).

To restore the redundancy configuration on the failed site after a non-catastrophic failure, do the following:

- 
- Step 1** As a root user, log into the failed site and unzip the RH\_ha.zip located on the “Disk 1 New Install” DVD. Unzipping **RH\_ha.zip** creates the /tmp/RH\_ha directory. Also, unzip the **RH\_ha.zip** in the primary location.
- Step 2** From the /tmp/RH\_ha directory, run the **perl resumeFromFailOver.pl -reconfigure\_setup** for restoring the replication between the failed and the active site.
- For information on options or arguments used with **resumeFromFailOver.pl** script, see [Table 6-1](#).
- ```
cd /tmp/RH_ha
perl resumeFromFailOver.pl -reconfigure_setup
```
- Step 3** Log into the active site and run the **perl resumeFromFailOver.pl -setup_replication** from the /tmp/RH_ha directory for setting up replication between the failed and the active site.
- For information on other options or arguments used with **resumeFromFailOver.pl** script, see [Table 6-1](#).
- Step 4** Confirm to continue with setting up of replication process.
- Step 5** To verify the setup, perform all of the tests (not including the step for creating database links) that are described in [Verifying the Geographical Redundancy Setup, page 4-14](#).
-

Stopping and Restarting Data Replication

These topics explain how to stop and restart the data replication process:

- [Stopping Data Replication, page 6-11](#)
- [Resuming Data Replication, page 6-13](#)

Stopping Data Replication

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 standby database. Then resume replication when maintenance is complete as described in [Resuming Data Replication, page 6-13](#).

The following includes an example of a stop replication session. In the example:

- The local active gateway is P1 with the IP address 10.33.150.66.
- The remote standby gateway is S1.

This example stops data replication from P1 to S1.



Note

This command *must* be run from the server that contains the standby database (S1 in this example). You can validate which server is the standby by running **primeha -status**.

- Step 1** Move to the correct directory.
- ```
cd /var/adm/cisco/prime-network/scripts/ha/util
```
- Step 2** Verify the status of the active and backup servers using **primeha -status**:
- ```
# primeha -status
```

```
+ Perl for HA already installed
```

```
HOST          Prime Network SERVICE          ORACLE SERVICE
S1            Standby Prime Network          Standby oracle   local
P1            Active Prime Network          Active oracle
```

Step 3 Log into the *server with the remote database* (S1) and enter the following command. This will stop replicating data and will shut down the remote site database.

```
# ./primeha -stop
```

Step 4 Enter the server details as shown in the following table.

Table 6-4 Prompts for primeha -stop and -start

Prompt for	Description
Remote server's gateway IP address	IP address for the primary gateway (P1). If the primary site has a local redundancy setup, enter the floating IP address for the Prime Network service. (Because the script is run from S1, P1 is the remote server.)
Root password for the node that has the gateway mounted	The root password for the operating system (required for SSH). (This would be the password for the P1 node.)
Remote database IP address	IP address for the primary database. If the primary site has a local redundancy setup, enter the floating IP address of the Oracle service. (Because the script is run from S1, P1 is the remote database.)
Root password for the node that has the database mounted	The root password for the operating system (required for SSH). (Again, this would be the password for the P1 node.)
Cluster sub server's IP address	(Displayed if standby database is part of a cluster) Enter the physical IP address of the remote node. (This would be the physical IP address for P1.)
Name for the OS user of the database	Example: oracle
Home directory of the user	Example: /opt/ora/oracle
Name for the OS user for Prime Network	Example: pn41
Home directory of the user Prime Network OS user	Example: /export/home/pn41

```
# ./primeha -stop
```

```
+ Perl for HA already installed
```

```
+ Stopping replication to remote node
```

```
- Enter the remote server's gateway IP address:
```

```
10.33.150.66
```

```
- Enter the root password for the node that has the gateway mounted:
```

```
pwd
```

```
- Enter the remote data base IP address:
```

```
10.33.150.66
```

```
- Enter the root password for the node that has the data base mounted:
```

```
pwd
```

```
- 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
pn41
- Enter the home directory of the user (pn41)
/export/home/pn41
- 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

Step 5 Verify the status of the active and backup servers using **primeha -status**:

```
# primeha -status
```

```
+ Perl for HA already installed
```

HOST	Prime Network SERVICE	ORACLE SERVICE
S1	Prime Network not running on this node	Oracle not running on this node
P1	Prime Network not running on this node	Oracle not running on this node

Step 6 Verify that all applications on the standby server (S1) are stopped by running **fsuser -c /export/home**. If any processes are still running (such as the Apache webserver), boot the standby server in single-user mode.

Step 7 Perform any necessary maintenance.

Resuming Data Replication



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 7 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 [Installing the Prime Network Gateway Geographical Redundancy Software, page 4-6](#) for information on using **setup_Prime_DR.pl** 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. Run this command after all work is completed on the remote site.

The following includes an example of a start replication session. In this example:

- The local active gateway P1 with IP address 10.33.150.66.
- The remote standby gateway is S1.

This example starts data replication from the local active gateway (P1) to the remote standby gateway (S1).



Note

This command must be run from the server that contains the remote database.

Step 1 Verify the following on the server with the remote database (S1):

- The database has not been down for more than seven days. If it has, you must recreate the remote database using **setup_Prime_DR.pl** script (see [Installing the Prime Network Gateway Geographical Redundancy Software](#), page 4-6).
- Maintenance on the server is complete.
- Whether the server is in single-user mode. If it is, reboot it in multi-user mode.

Step 2 Log into the server with the remote database (S1) and move to the correct directory.

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

Step 3 Verify the status of the active and backup servers using **primeha -status**. If any services are running, stop them using **primeha -stop**.

Step 4 Enter the following command. This will start replicating data and will shut down the remote site database.

```
# ./primeha -start
```

Step 5 Enter the server details (see [Table 6-4](#)).



Note If the process aborts, run **primeha -stop** again. (The script most likely aborted because a process is not shut down.) Then verify that no services are running with **primeha -status**.

Step 6 When the process completes, verify the status of the active and backup servers using **primeha -status**:

```
# primeha -status
```

```
+ Perl for HA already installed
```

HOST	Prime Network SERVICE	ORACLE SERVICE	
S1	Standby Prime Network	Standby oracle	local
P1	Active Prime Network	Active oracle	

Changing IP Addresses after a Failover or Switchover

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 6-14
- [Changing the Gateway IP Address on a Single Unit \(switchUnit.pl\)](#), page 6-16

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 or arguments 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* in the [Cisco Prime Network 4.3 Administrator Guide](#).

Option/Arguments	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 <i>new-gateway-ip</i>	IP address of the gateway that is running after the failover or switchover.
-newdbip <i>new-database-ip</i>	IP address of the database that running after the failover or switchover.
-oldgwip <i>old-gateway-ip</i>	IP address of the gateway that was running prior to the failover or switchover.
-olddbip <i>old-database-ip</i>	IP address of the database that was running prior to the failover or switchover.
[-newldapurl <i>new-ldap-url</i> -oldldapurl <i>old-ldap-url</i>]	(LDAP only) URL for the LDAP server that will be used by the running gateway (<i>new-ldap-url</i>), and the URL that was used by the gateway that was running prior to the failover or switchover (<i>old-ldap-url</i>). Use the following format: ldap://host.company.com:port
[-newldapprefix <i>new-ldap-prefix</i> -oldldapprefix <i>old-ldap-prefix</i>]	(LDAP only) First part of the LDAP DN (which is used to uniquely identify users) for the new and old LDAP server. Both <i>new-ldap-prefix</i> and <i>old-ldap-prefix</i> 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 <i>new-ldap-suffix</i> -oldldapsuffix <i>old-ldap-suffix</i>]	(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 <i>new-ldap-suffix</i> and <i>old-ldap-suffix</i> should use the following format ,CN=Users,DC=LDAP_server,DC=company,DC=com
[-newldapisssl <i>new-ldap-is-ssl</i> -oldldapisssl <i>old-ldap-is-ssl</i>]	(LDAP only) Encryption protocol to be used for communication between the running Prime Network gateway server and the new LDAP server (<i>new-ldap-is-ssl</i>), and the protocol that was used between the old gateway and LDAP servers (<i>old-ldap-is-ssl</i>).

Step 1 If you will reset LDAP information, reconfigure them first from the Prime Network Administration GUI client. For details on LDAP, see *Configuring Prime Network to Communicate with the External LDAP Server* in the [Cisco Prime Network 4.3 Administrator Guide](#).

Step 2 Log into the primary gateway server as *pnuser*.

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]
[-newldapissql new-ldap-is-ssl -oldldapissql 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 1.1.1.1.
- The site was switched over to the standby gateway (1.1.1.2) and database (1.1.1.3).

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 1.1.1.2 for the gateway and 1.1.1.3 for the database. In this example the system is not using LDAP, so those parameters are not included.

```
hostname% cd $ANAHOME/Main/ha
hostname% perl changeSite.pl -force -newgwip 1.1.1.2 -newdbip 1.1.1.3 -oldgwip 1.1.1.1
-olddbip 1.1.1.1

Thu Apr 14 16:08:22 2013 --[INFO]: '-Forced change of gw address from 1.1.1.1 to
1.1.1.2.... '
Thu Apr 14 16:08:22 2013 --[INFO]: '--changing uplinks for gw AVM0'
Thu Apr 14 16:08:22 2013 --[INFO]: '--changing uplinks for unit AVM0s'
Thu Apr 14 16:08:22 2013 --[INFO]: '--changing gw ip and haservice for unit AVM99s'
Thu Apr 14 16:08:22 2013 --[INFO]: '--changing registry on units'
Thu Apr 14 16:08:37 2013 --[INFO]: '--changing localhost entry for gw AVM99'
Thu Apr 14 16:08:37 2013 --[INFO]: '-Forced change of db server address from 1.1.1.1 to
1.1.1.3.... '
Thu Apr 14 16:08:37 2013 --[INFO]: '--changing db server ip for gw AVM66'
Thu Apr 14 16:08:38 2013 --[INFO]: '--changing db server ip for gw persistency.xml'
Thu Apr 14 16:08:38 2013 --[INFO]: '--changing db server ip for template persistency.xml'
Thu Apr 14 16:08:38 2013 --[INFO]: '--changing db server ip for unit persistency.xml'
Thu Apr 14 16:08:38 2013 --[INFO]: '-Forced change of NCCM address from 1.1.1.1 to
1.1.1.3.... '
new IP address is: 1.1.1.3
jdbc.properties file has been updated to change to new IP address
Thu Apr 14 16:08:39 2013 --[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 *pnuser*.

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
```
