



Functional Overview of Gateway Local Redundancy and Geographical Redundancy

These topics provide a functional overview of the Prime Network gateway high availability solutions:

- [Local Redundancy Functional Overview, page 2-1](#)
- [Geographical Redundancy Functional Overview, page 2-5](#)



Note

Gateway high availability is supported only when the gateway software, Oracle database, and Infobright database (applicable for Operations Report) are installed on the same server.

Local Redundancy Functional Overview


Gateway local redundancy configuration uses one active and one standby node to provide an automatic failover solution for local hardware faults, without the need to reconfigure IP addresses. The solution uses the Red Hat Cluster Suite (RHCS). The nodes are monitored by RHCS and if the node managing the services fails, the services are seamlessly moved to the other node. In case of a single service failure, the cluster will attempt to restart the service. If the retries fail, the service will be relocated to the second node and started on that node. This does not impact the other service in the cluster.

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

[Figure 2-1](#) shows a basic dual-node cluster local redundancy configuration, where the Prime Network gateway service is on Server P1, the Oracle database service is on Server P2, and both servers are connected to a server for external storage. Both servers use an embedded database.

The RHCS local redundancy solution requires a fencing device, which is a hardware unit that disconnects a node from shared storage to ensure data integrity. For more information on fencing options, see [Fencing Options, page 2-3](#).

Figure 2-1 Architecture for Gateway Local Redundancy



Configuration Details for Local Redundancy

Local redundancy requires that RHCS be installed on both nodes. Out of the box, both services run on the node from which the installation script is run. This configuration can be changed, if desired, using RHCS web GUI or CLI (**clusvcadm** utility). For details on the required system configuration for local redundancy, see [Installation Requirements for Local Redundancy, page 3-4](#)

The local redundancy setup has the following:

- [Dual Node Cluster, page 2-2](#)
- [RHCS Installed on Both Nodes, page 2-2](#)
- [External Shared Storage, page 2-3](#)
- [Fencing Options, page 2-3](#)
- [Security, page 2-5](#)

Dual Node Cluster

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

RHCS Installed on Both Nodes

RHCS manages the local redundancy by monitoring cluster configured services: **ana** and **oracle_db**. If you have installed Operations Reports, RHCS also manages the Infobright database service, **ifb**. If a hardware or software failure occurs, the RHCS automatically restarts the failed node's services on the functional node.



Note

Operations Reports are only available to customers with Operations Reports license prior to May 2018. For re-installation of Operations Reports contact a Cisco customer support representative.

Table 2-1 lists the services that are monitored by RHCS.

Table 2-1 Cluster Configured Services Monitored by RHCS

RHCS Service	Description	
ana	Monitors Prime Network (AVM 99), the Prime Network Integration Layer (if installed), and Prime Network Operations Reports (if installed). It consists of the following resources.	
	IP address	<i>ana_service_floating_IP</i>
	Scripts	ana.sh (Prime Network) pcil.sh (Prime Network Integration Layer)
oracle_db	Monitors Oracle processes and listener and consists of the following resources.	
	IP address	<i>oracle_db_floating_IP</i>
	Scripts	oracles.sh, lsnr.sh
ifb	Monitors Infobright processes and XXX and consists of the following resources:	
	IP address	<i>infobright_db_floating_IP</i>
	Scripts	infobright.sh

The floating IP address is different from either node's physical IP address. The floating IP address is associated with a service rather than a particular machine in the cluster. Therefore, the *cluster IP address* is the *floating IP address of the management port of the cluster*. It floats because it always points the parent device (for example, it would change from P1 to P2 in case of a switchover or failover).

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 and the Infobright service groups, then the Prime Network service group. (RHCS does not enforce this behavior.)

External Shared Storage

RHCS requires an external shared storage that is mountable from both nodes. The external shared storage contain the Prime Network, Oracle, and (if Operations Reports is installed) Infobright files.



Note

Operations Reports are only available to customers with Operations Reports license prior to May 2018. For re-installation of Operations Reports contact a Cisco customer support representative.

Fencing Options

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

The supported fencing options are:

- **fence_ipmilan**—Intelligent Platform Management Interface (IPMI) v1.5 or higher compliant devices over a LAN.
- **fence_ilo**—Hewlett Packard Integrated Lights Out (HP iLO).
- **fence_vmware_soap**—VMware with SOAP API. This agent communicates with the VMware vCenter server that is managing the VM that will be fenced. If you choose this fencing method, you will be prompted for the vCenter user login and password, the vCenter IP address, and the vCenter hostname. Not all RHEL versions support this option; see [Hardware and Software Requirements for Local Redundancy, page 3-5](#) and the [Red Hat site](#).
- **fence_manual**—This option allows you to add a Red Hat-supported fencing device not listed above. If you choose Manual, the **fence-manual-fencing** agent is assigned. This fencing agent is temporary and should not be used in production because it does not perform automated actions. If a cluster problem occurs, the node and storage must be manually disconnected, or another fencing agent must be used to disconnect them. If you choose this option during the installation because you want to add a different Red Hat-supported fencing device, provision the device *after* installation using the RHCS GUI. When you add it, be sure to add it as the main fencing method, and move the manual fencing agent to the backup method, as shown in [Figure 2-2](#).

General information about the RHCS web GUI is provided in [Configuring the RHCS Web Interface \(Optional\), page 3-24](#). However, see the Red Hat Conga documentation for complete information about using the RHCS web GUI application. Additionally, you need the RHCS user documentation to provision and manage cluster fencing devices. See the [Red Hat site](#) for more information.

**Note**

Keep these items in mind:

- To prevent fencing loops, the cluster interconnect and power fencing (for example, HP-iLO) should use the same network, such as bond0.
- If the main fencing device is a remote power switch, define all ports that supply power to the node simultaneously.
- If manual fencing is used, before disconnecting the node, remove the cman and rgmanager services from the automatic startup sequence. For more information on the command to remove these services, see [Manually Fencing, page 3-28](#).

Figure 2-2 RHCS GUI Fencing Method Window

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

Security

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

Geographical Redundancy Functional Overview

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

Geographical redundancy can be installed alone (geographical redundancy *only*) or with local redundancy (local + geographical redundancy), depending on your configuration; see [Table 1-2 on page 1-3](#). In both cases, the remote site (S1) contains its own server, database, and storage—all located at another geographical location. The DR node at the remote site is the backup to the node primary site. [Figure 2-3](#) illustrates a deployment with geographical redundancy *alone*. It includes a single node with external storage in the primary site (P1), and a single node with external storage in the remote site (S1).

Figure 2-3 Architecture for Gateway with Geographical Redundancy Alone





Figure 2-4 illustrates a deployment with both local + geographical redundancy. It includes a dual-node cluster with external shared storage in the primary site (P1 and P2), and a single node with external storage in the remote site (S1).

Figure 2-4 Architecture for Gateway with Local Redundancy and Geographical Redundancy

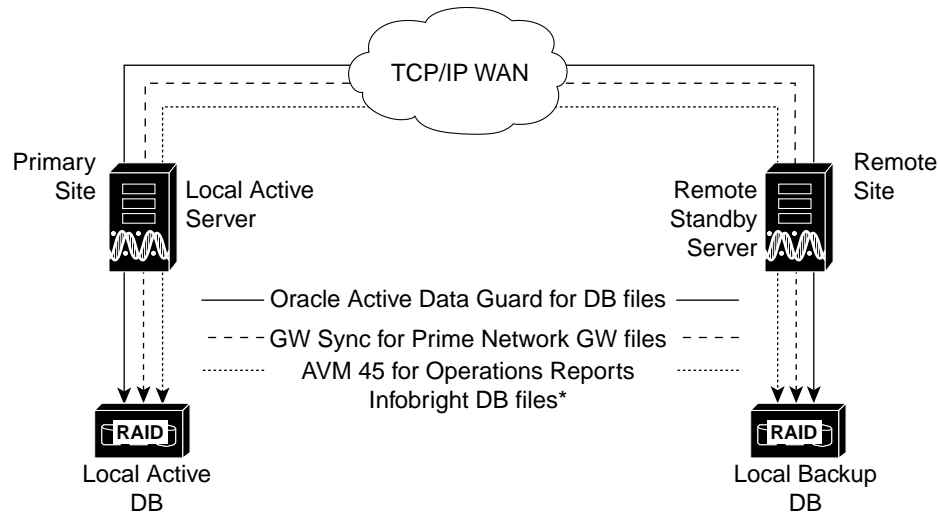
**Note**

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 site (or vice versa).

The data stored in the server and Oracle database is continuously replicated between the two sites. The primary and standby Oracle database are monitored and synchronized using ADG. If Operations Reports is installed, AVM 45 performs the synchronization between the two Infobright databases and Oracle database.

The Prime Network server files (registry and system files) are synchronized using the GWSync utility. Prime Network periodically monitors and validates the replication process and issues a System event in case of a problem. Figure 2-5 shows the data replication process between the primary site and remote site. To secure the channel used for data replication, an SSH key exchange is performed during the Prime Network installation.

Figure 2-5 Replication Configuration for Geographical Redundancy



* Gateway, Oracle DB, and Infobright DB are on same server

For disaster recovery (if the primary site becomes unavailable), a manual failover can be triggered from the remote site. The utilities for managing the manual failover are described in [Maintaining Geographical Redundancy](#), page 4-16.

The geographical redundancy solution uses the following replication processes.

- [Oracle ADG Replication Process](#), page 2-7
- [Gateway Sync \(GWSync\) Replication Process](#), page 2-8
- [Infobright Database Replication Process \(Operations Reports\)](#), page 2-9

Oracle ADG Replication Process

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

When the high availability solution is installed, it sets up the cron jobs that will monitor the synchronization process.

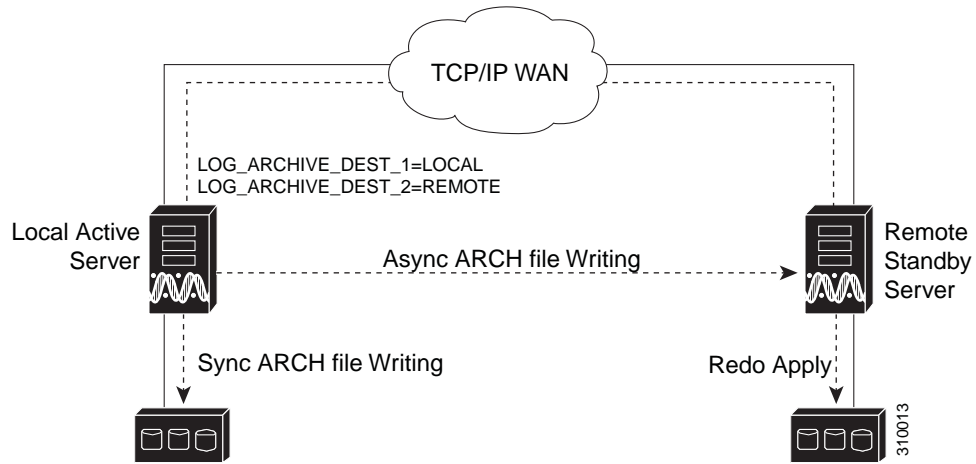


Note

ADG uses port 1521 for communication between the servers. This port must be open.

[Figure 2-6](#) illustrates how data is replicated between the primary site database and the remote site.

Figure 2-6 ADG Database Replication Process (ADG Geographical Redundancy)



Note

The databases must have identical disk capacities and mount points.

To troubleshoot problems with the replication process, see the [Verifying the Geographical Redundancy Setup](#), page 4-14.



Note

When the `emdbctl --restore` command is used with Oracle ADG, reconfigure the Oracle database replication after restoring the primary database.

Gateway Sync (GWSync) Replication Process

Gateway Sync (GWSync) is a RHEL rsync utility that replicates the Prime Network home directory (and any file system data that is required for disaster recovery) from the primary gateway to the remote site. The GWSync process replicates the gateway data between servers. If Operations Reports is installed, GWSync replicates the folders, system files, and registry data used by Operations Reports.



Note


Operations Reports are only available to customers with Operations Reports license prior to May 2018. For re-installation of Operations Reports contact a Cisco customer support representative.

Cron jobs trigger synchronization at both the primary and remote sites. Data is exchanged using SSH across secure channels. GWSync only sends data that has changed.

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

Prime Network monitors and validates the replication processes (ADG and Gateway Sync) and issues a system event if replication problems occur. To troubleshoot problems with the replication process, see [Verifying the Geographical Redundancy Setup](#), page 4-14.

Figure 2-7 How GWSync Replication Process is Monitored (ADG Geographical Redundancy)



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

Infobright Database Replication Process (Operations Reports)

The Infobright database is the repository for the event data used by Operations Reports. This primary Infobright database is synchronized with the remote Infobright database using AVM 45. The Infobright backup files are created on an hourly basis, and only after the file is closed it is loaded to the remote Infobright database using AVM45. In case of a failover, archived event data and standard event data from the previous 80 minutes is lost. This is due to the time required to load and validate the data that is saved on the secondary Infobright database.



Note

Operations Reports are only available to customers with Operations Reports license prior to May 2018. For re-installation of Operations Reports contact a Cisco customer support representative.

ESXi Testing Details

The nodes run on ESXi 5.0 on the Cisco UCS Blade Server.

