# Setting Up RDU Redundancy

The RDU (Regional Distribution Unit) redundancy feature involves setting up the RDU in High Availability (HA) mode where a two node failover pair is configured for the RDU.

This chapter describes how to configure the RDU redundancy feature in Prime Cable Provisioning.

## Prerequisites

The following prerequisites must be met before you proceed with RDU redundancy setup:

**Utility Requirements**

- Identify the two servers on which you wish to configure the primary and secondary RDU nodes.

**Redundancy Requirements**

- Redundant network configuration should be available to avoid network downtime.
- Redundant electrical supply must be available on both servers. Ensure that the electrical supply source for both servers is reachable.

**PCP Geo Redundancy Requirements**

Route injection for VIP (virtual IP) needs to be done on the ingress routers to which primary and secondary servers are connected.

The VIP will be advertised as RIP2 advertisement from the active server ,so route redistribution needs to be done for RIP2 to the dynamic routing protocol running in the user environment.

***Example:*** Here OSPF is the dynamic protocol

router ospf 1

redistribute rip metric-type 1 subnets.

### Logical Volume Manager (LVM) Setup

Both RDU nodes must be configured over Logical Volume Manager (LVM). The LVM allows you to create a volume group which can be further divided into logical volumes based on the requirement. The LVM also provides the flexibility to resize the volume group and logical volumes based on the dynamic memory usage.

The LVM setup involves the following considerations:

1. On both primary and secondary RDU nodes, a logical volume group must be created with three logical volumes on it. The logical volumes are created based on the following specifications:
    - <*logical volume for Prime Cable Provisioning install directory*> - Mounted on /bprHome directory. For example, LVBPRHOME.
    - <*logical volume for Prime Cable Provisioning data directory*> - Mounted on /bprData directory. For example, LVBPRDATA
    - <*logical volume for Prime Cable Provisioning log directory* > - Mounted on /bprLog directory. For example, LVBPRDBLOG
2. Ensure that the /etc/fstab entries are updated for these logical volumes.
3. Ensure that the /bprData, /bprHome, and /bprLog directories are empty.
4. The logical volumes should be of same capacity on both the nodes with a pre-created ext4 filesystem.

### Requirements for Proper Synchronization between Nodes

- Identify the virtual IP address that can be used to locate both servers. Ensure that this virtual IP address is not configured for any physical interfaces in the network infrastructure. Also, this virtual IP address must belong to the network cluster in which you configure the primary and secondary RDU nodes.
- The Network Time Protocol (NTP) must be synchronized between primary and secondary RDU nodes.
- Both RDU nodes must exist in the same network cluster. This ensures that the same virtual IP address can be used to reach both the RDU nodes.
- Both RDU nodes must be configured with the following two network interfaces:
    - Public access - Used for external communication. The DPE, PWS, CPNR_EP, and API clients use this interface to communicate with the RDU. A public IP address is configured to access this network interface.
    - Failover link between RDU nodes - Used for disk synchronization. If the RDU nodes are co-located, the failover link can be a crossover link, else the failover link must be configured in the private LAN. The failover IP addresses for primary and secondary nodes must be unique in the network cluster. This is to achieve the high speed access between primary and secondary RDU nodes, and also make them isolated from the network traffic.

### Miscellaneous Requirements

- Both RDU nodes must be configured with Gigabyte network interfaces.

The following figure provides a high level RDU HA setup.

*Figure 1: RDU HA Setup*

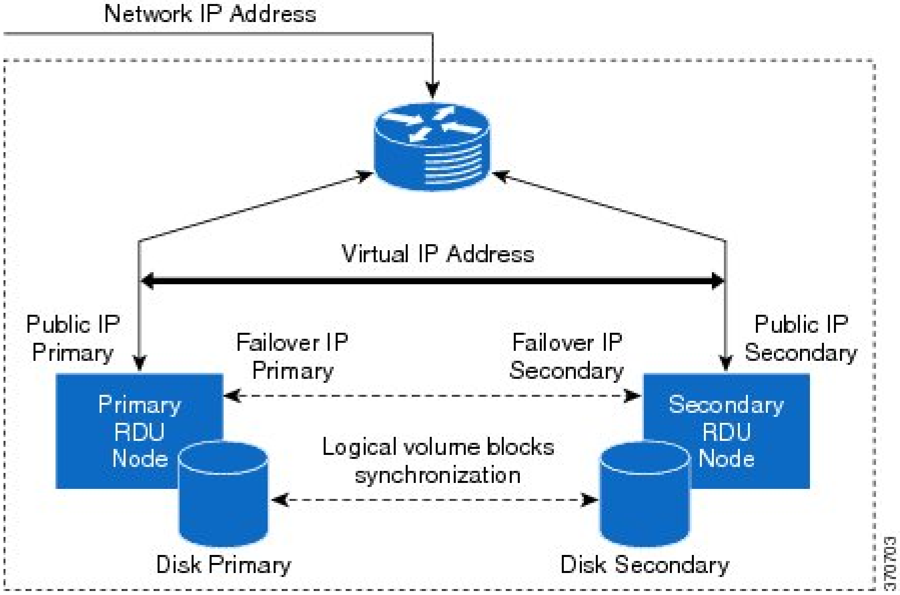

# RDU HA Setup in Primary-Secondary Mode

If both primary and secondary RDU nodes are available in the network infrastructure, you can use primary-secondary installation mode to configure HA cluster.

To configure the RDU HA setup using primary-secondary installation mode, you must perform the following tasks sequentially:

1. Configure RDU nodes for HA setup
2. Set up RDU two node failover pair

The following figure describes the workflow for configuring RDU HA setup using primary-secondary installation mode.

*Figure 2: RDU HA Setup - primary-secondary mode*

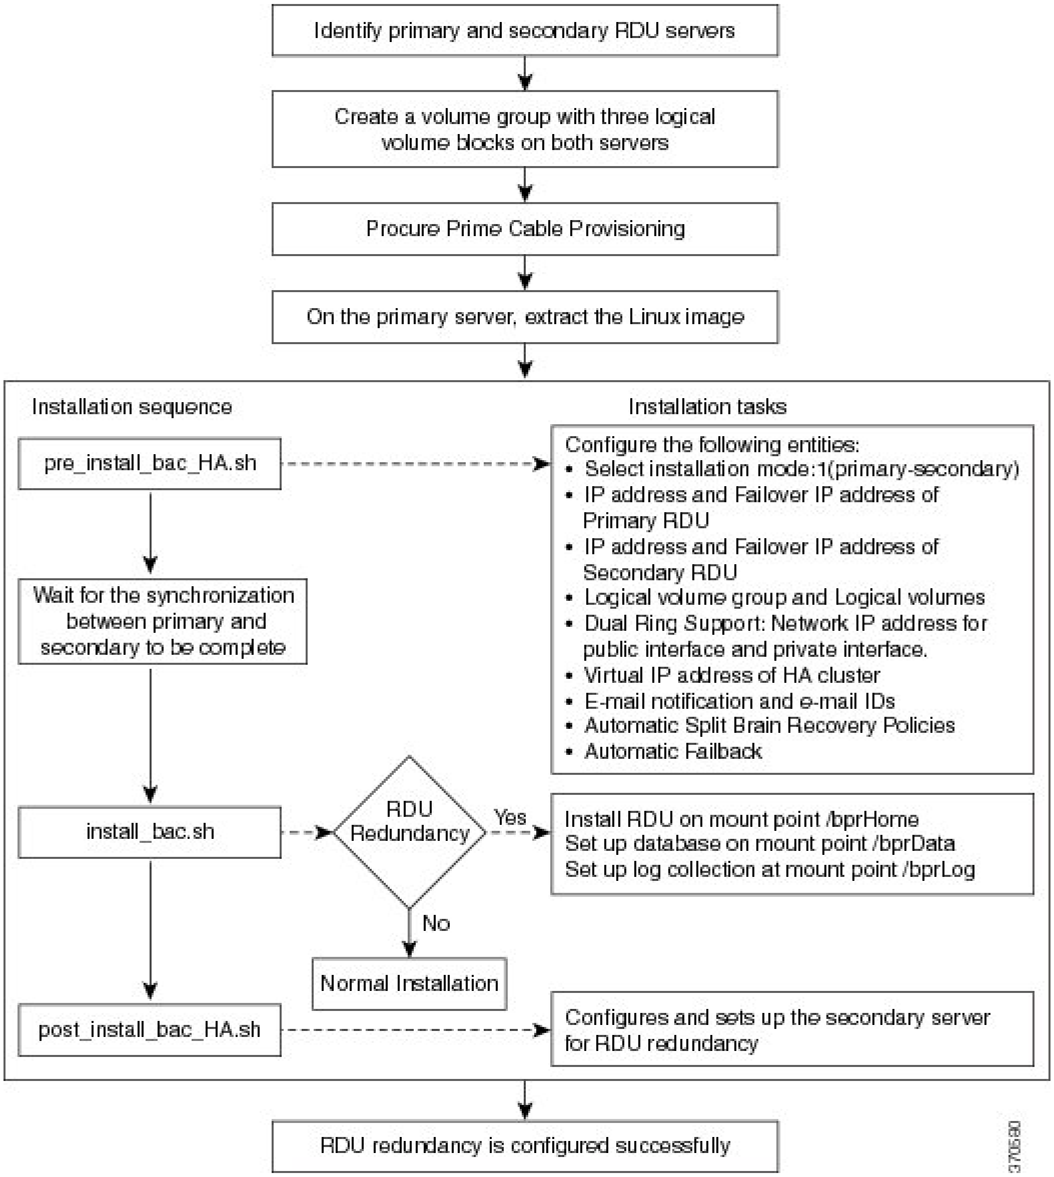
Identify primary and secondary RDU servers
Create a volume group with three logical volume blocks on both servers
Procure Prime Cable Provisioning
On the primary server, extract the Linux image
Installation sequence
Installation tasks
pre_install_bac_HA.sh
Configure the following entities:
• Select installation mode:1(primary-secondary)
• IP address and Failover IP address of Primary RDU
• IP address and Failover IP address of Secondary RDU
• Logical volume group and Logical volumes
• Dual Ring Support: Network IP address for public interface and private interface.
• Virtual IP address of HA cluster
• E-mail notification and e-mail IDs
• Automatic Split Brain Recovery Policies
• Automatic Failback
Wait for the synchronization between primary and secondary to be complete
install_bac.sh
RDU Redundancy
Yes
Install RDU on mount point /bprHome
Set up database on mount point /bprData
Set up log collection at mount point /bprLog
No
Normal Installation
post_install_bac_HA.sh
Configures and sets up the secondary server for RDU redundancy
RDU redundancy is configured successfully
370590

## Preparing RDU Nodes for HA Setup in Primary-Secondary Mode

Before installing the RDU redundancy function, you must perform the required server configurations and establish a communication channel between primary and secondary RDU servers.

To prepare primary and secondary RDU nodes for HA setup:

### Procedure

**Step 1** Configure the SSH access between primary and secondary RDU nodes and disable the password authentication:

1. On the primary RDU node, run the following command to create the SSH keys:

```
# ssh-keygen -t dsa -f ~/.ssh/id_dsa -N ""
```

The system generates the SSH keys as follows:

```
Generating public/private dsa key pair.
/root/.ssh/id_dsa already exists.
Overwrite (y/n)? y
Your identification has been saved in /root/.ssh/id_dsa.
Your public key has been saved in /root/.ssh/id_dsa.pub.
The key fingerprint is:
14:50:6f:46:39:67:35:89:37:87:55:29:90:fd:49:0b root@<rdu-primary>
The key's randomart image is:
+--[ DSA 1024]----+
| .oo .o+oo+=|
| +o.+E=+.|
| . ++ .=oo|
| . o + |
| S |
| |
| |
| |
| |
+-----------------+
```

2. Copy the SSH keys to the secondary RDU node:

```
# cp ~/.ssh/id_dsa.pub ~/.ssh/authorized_keys
# scp -r ~/.ssh <rdu-secondary>:
```

The system prompts you to enter the password for secondary RDU node. Once you enter the password, the system copies the SSH keys to the secondary RDU node:

```
root@<rdu-secondary>'s password:
known_hosts
100% 1199 1.2KB/s 00:00
id_dsa.pub
100% 606 0.6KB/s 00:00
authorized_keys
100% 606 0.6KB/s 00:00
id_dsa
100% 668 0.7KB/s 00:00
```

3. On the primary and secondary RDU servers, verify if the SSH access is available without password authentication:

```
# ssh <target RDU server> -- uname -n
```

where, <target RDU server> is the server name of primary or secondary RDU. For example, <rdu-primary> or <rdu-secondary>. If you verify the SSH access from primary RDU, you must enter the secondary RDU as the <target RDU server>.

The following output confirms that SSH access is enabled between primary and secondary RDU nodes:

```
Linux <target RDU server> 2.6.32-279.e16.x86_64 #1 SNP Wed Jun 13
18:24:36 EDT 2012
 x86_64 x86_64 x86_64 GNU/Linux
```

**Step 2** Verify if the network access is available and the ethernet links are functioning normal:

```
# ifconfig
```

The following output confirms that the network access is available:

```
eth0 Link encap:Ethernet HWaddr 00:0C:29:ED:AE:75
          inet addr:10.104.182.92 Bcast:10.104.182.255 Mask:255.255.255.0

            inet6 addr: fe80::20c:29ff:feed:ae75/64 Scope:Link
            UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
            RX packets:761 errors:0 dropped:0 overruns:0 frame:0
            TX packets:272 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:1000
            RX bytes:66998 (65.4 KiB) TX bytes:41753 (40.7 KiB)

eth1        Link encap:Ethernet HWaddr 00:0C:29:ED:AE:7F
            inet addr:192.168.1.1 Bcast:192.168.1.255 Mask:255.255.255.0
            inet6 addr: fe80::20c:29ff:feed:ae7f/64 Scope:Link
            UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
            RX packets:6 errors:0 dropped:0 overruns:0 frame:0
            TX packets:13 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:1000
            RX bytes:360 (360.0 b) TX bytes:830 (830.0 b)
lo          Link encap:Local Loopback
            inet addr:127.0.0.1 Mask:255.0.0.0
            inet6 addr: ::1/128 Scope:Host
            UP LOOPBACK RUNNING MTU:16436 Metric:1
            RX packets:0 errors:0 dropped:0 overruns:0 frame:0
            TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:0
            RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

**Step 3** Verify the accessibility between primary and secondary RDU nodes:

- Ping secondary RDU node from the primary server:

  ```
  # ping <rdu-secondary>
  ```

- Ping primary RDU node from the secondary server:

  ```
  # ping <rdu-primary>
  ```

# Setting Up RDU Two Node Failover Pair

The RDU two node failover pair setup involves running three installation scripts sequentially on the primary RDU. The following table describes each installation script and the sequence in which you run them.

*Table 1: RDU Two Node Failover Pair Setup*

| Sequence | Installation Scripts | Description |
|---|---|---|
| 1 | pre_install_bac_HA.sh | Used to configure the following entities:<br>• Logical volume group and logical volumes<br>• Failover IP addresses of both the RDU nodes<br>• Network IP address<br>• Automatic failback<br>• Virtual IP address<br>• Secondary RDU server public IP address<br>• Utilities required for RDU redundancy function setup. |
| 2 | install_bac.sh | Used to install the RDU component on the logical volumes. |
| 3 | post_install_bac_HA.sh | Used to automate the supported configuration tasks for RDU redundancy function setup. |

To set up the RDU two node failover pair:

### Procedure

**Step 1** Log into the primary RDU as root.

**Step 2** Extract the installation package using the following commands:

**Step 3** Run the preinstallation script using the following command:

```
# sh pre_install_bac_HA.sh
```

**Step 4** Enter the RDU redundancy information. The RDU redundancy information includes:

- Failover IP addresses of primary and secondary RDU servers - Unique IP address in the network cluster that are used for data synchronization between primary and secondary RDU nodes.
- Public IP address of the secondary RDU - Public IP address that is used for external communication.

**Step 5** Enter the name of the logical volume group. For example, VGBPR.

**Step 6** Enter the name of the logical volumes created for home, data, and database log directories. For example, you can enter **LVBPRHOME** for home directory, **LVBPRDATA** for data directory, and **LVBPRDBLOG** for database log directory.

**Step 7** Enter the network IP address of the subnet under which both the RDU cluster nodes exist.

**Step 8** Enter the virtual IP address. The virtual IP address is the floating IP address that is used to reach both primary and secondary RDU nodes, but allocated to only active RDU node. Ensure that this virtual IP address is not configured for any physical interfaces in the network.

**Step 9** Enter **y** to enable automatic failback, else enter **n**. If the automatic failback is enabled, the primary RDU node becomes active once it comes up after the failover event.
The preinstallation script performs the following automated operations on both primary and secondary RDU nodes:

- Installs the utilities for RDU redundancy setup.
- Creates the required resources for synchronizing the block devices.

**Step 10** Install RDU on the synchronized logical volumes; **LVBPRHOME**, **LVBPRDATA**, and **LVBPRDBLOG**. For details, see Installing the RDU in Interactive Mode.

**Step 11** Run the post-installation script available under directory:

```
# sh post_install_bac_HA.sh
```

The post-installation script performs the automated configuration tasks required for RDU redundancy function setup.

# RDU HA Uninstall Scripts

To uninstall the RDU HA setup in the primary and secondary RDU nodes, you must run the uninstall scripts sequentially. The following table describes the uninstall scripts and the sequence in which you run them.

*Table 2: Uninstall RDU HA Setup*

| Sequence | Uninstall Scripts | Description |
|---|---|---|
| 1 | uninstall_bac.sh | Used to unistall the RDU component from the logical volume blocks; /bacHAHome, /bacHAData, and /bacHADBLog . For details, see Uninstalling Prime Cable Provisioning. |
| 2 | uninstall_bac_HA.sh | Used to remove the utilities installed for RDU redundancy function. |

# RDU Geo Redundancy

RDU Geo Redundancy is an enhanced feature of RDU HA supported on RHEL 6.5 or CentOs 6.5 (both 64bit), wherein the RDU primary and secondary node can be in different geographical location or both the nodes can be in different subnet.

- In Geo redundancy mode the VIP can be in any subnet it is not necessary to have in the subnet range common to both nodes.
- In Geo redundancy mode the CIDR value of VIP should be 32.
- The VIP will be advertised as a RIP advertisement from the active server, so on the ingress router of both the nodes route injection need to be done.

For setting up RDU in Geo redundancy mode follow steps mentioned in Setting Up RDU Redundancy, on page 1