

Prime Cable Provisioning 6.1.5 RDU High Availability Deployment with Geo Mode Redundancy

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components](#)

[Installation](#)

[Network Diagram](#)

[1. LVM Create Volume for LVBPRHOME, LVBPRDATA and LVBPRDBLOG on Both Servers](#)

[2. Prepare Linux 7.4 server for RDU HA Deployment on Both Servers](#)

[3. Install RDU server in Geo-Redundancy Mode](#)

[4. Layer-3 Routing Prerequisite for Geo-Redundancy Deployment](#)

[RDU Geo-Redundancy](#)

[PCP Geo-Redundancy Requirements](#)

[Post-Check HA](#)

Introduction

This document describes the installation of Prime Cable Provisioning 6.1.5 in High Availability (HA) with geo mode redundancy.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Redhat Linux knowledge and understanding of file system and partitioning.
- Install 6.1.5 RHEL 7.4/Kernel 3.10.0-693.11.6.x86_64 on new primary and secondary virtual/physical machine. RDU HA with geo mode is only compatible with this RHEL OS and kernel version and its rpm packages.
- Knowledge of Linux DRBD file storage replication method and Corosync-pacemaker cluster concept.
- The network configuration file should contain only the system hostname, not the Fully Qualified Domain Name (FQDN).

Components

The information in this document is based on these software and hardware versions:


- Platform: Red Hat Linux 7.4
- Software: Prime Cable provisioning 6.1.5 image.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Installation

Network Diagram

Figure 1. RDU Redundancy



1. LVM create volume for LVBPRHOME, LVBPRDATA and LVBPRDBLOG on both servers.
2. Preparing Linux 7.4 server for RDU HA deployment on both servers.
3. Installation of RDU server in Geo redundancy mode
 - Installation of RDU server in Geo redundancy mode.
 - Pre check HA. -RDU HA Setup in Primary- Secondary Mode.
 - Install HA. - Install 6.1.5 PCP instance.
 - Post check HA.
4. Layer-3 Routing prerequisite for Geo redundancy deployment.

1. LVM Create Volume for LVBPRHOME, LVBPRDATA and LVBPRDBLOG on Both Servers

This illustration is done for the Secondary server. The same procedure needs to be completed on the primary server as well.

- Add new partition as sda3 and allocate disk with the use of fdisk command.

```
[root@pcprdusecondary ~]# fdisk -l
```

```
Disk /dev/sda: 107.4 GB, 107374182400 bytes, 209715200 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```

```
Sector size (logical/physical): 512 bytes / 512 bytes
```

```
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk label type: dos
```

```
Disk identifier: 0x00025a26
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	2048	2099199	1048576	83	Linux
/dev/sda2		2099200	31211519	14556160	8e	Linux LVM

```
Disk /dev/mapper/rhel-root: 4294 MB, 4294967296 bytes, 8388608 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```

```
Sector size (logical/physical): 512 bytes / 512 bytes
```

```
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk /dev/mapper/rhel-swap: 8455 MB, 8455716864 bytes, 16515072 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```

```
Sector size (logical/physical): 512 bytes / 512 bytes
```

```
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk /dev/mapper/rhel-home: 2147 MB, 2147483648 bytes, 4194304 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```

```
Sector size (logical/physical): 512 bytes / 512 bytes
```

```
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
[root@pcprdusecondary ~]# fdisk /dev/sda
```

```
Welcome to fdisk (util-linux 2.23.2).
```

Changes will remain in memory only, until you decide to write them.

Be careful before using the write command.

Command (m for help): **m**

Command action

- a toggle a bootable flag
- b edit bsd disklabel
- c toggle the dos compatibility flag
- d delete a partition
- g create a new empty GPT partition table
- G create an IRIX (SGI) partition table
- l list known partition types
- m print this menu
- n add a new partition
- o create a new empty DOS partition table
- p print the partition table
- q quit without saving changes
- s create a new empty Sun disklabel
- t change a partition's system id
- u change display/entry units
- v verify the partition table
- w write table to disk and exit
- x extra functionality (experts only)

Command (m for help): **p**

Disk /dev/sda: 107.4 GB, 107374182400 bytes, 209715200 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk label type: dos

Disk identifier: 0x00025a26

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	2048	2099199	1048576	83	Linux
/dev/sda2		2099200	31211519	14556160	8e	Linux LVM

Command (m for help): **n**

Partition type:

p primary (2 primary, 0 extended, 2 free)
e extended

Select (default p): **p**

Partition number (3,4, default 3): **3**

First sector (31211520-209715199, default 31211520):

Using default value 31211520

Last sector, +sectors or +size{K,M,G} (31211520-209715199, default 209715199):

Using default value 209715199

Partition 3 of type Linux and of size 85.1 GiB is set

Command (m for help): **p**

Disk /dev/sda: 107.4 GB, 107374182400 bytes, 209715200 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk label type: dos

Disk identifier: 0x00025a26

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	2048	2099199	1048576	83	Linux
/dev/sda2		2099200	31211519	14556160	8e	Linux LVM
/dev/sda3		31211520	209715199	89251840	83	Linux

Command (m for help): **t**

Partition number (1-3, default 3): **3**

Hex code (type L to list all codes): **L**

0	Empty	24	NEC DOS	81	Minix / old Lin	bf	Solaris
---	-------	----	---------	----	-----------------	----	---------

1	FAT12	27	Hidden NTFS Win	82	Linux swap / So	c1	DRDOS/sec (FAT-
2	XENIX root	39	Plan 9	83	Linux	c4	DRDOS/sec (FAT-
3	XENIX usr	3c	PartitionMagic	84	OS/2 hidden C:	c6	DRDOS/sec (FAT-
4	FAT16 <32M	40	Venix 80286	85	Linux extended	c7	Syrinx
5	Extended	41	PPC PReP Boot	86	NTFS volume set	da	Non-FS data
6	FAT16	42	SFS	87	NTFS volume set	db	CP/M / CTOS / .
7	HPFS/NTFS/exFAT	4d	QNX4.x	88	Linux plaintext	de	Dell Utility
8	AIX	4e	QNX4.x 2nd part	8e	Linux LVM	df	BootIt
9	AIX bootable	4f	QNX4.x 3rd part	93	Amoeba	e1	DOS access
a	OS/2 Boot Manag	50	OnTrack DM	94	Amoeba BBT	e3	DOS R/O
b	W95 FAT32	51	OnTrack DM6 Aux	9f	BSD/OS	e4	SpeedStor
c	W95 FAT32 (LBA)	52	CP/M	a0	IBM Thinkpad	hi	BeOS fs
e	W95 FAT16 (LBA)	53	OnTrack DM6 Aux	a5	FreeBSD	ee	GPT
f	W95 Ext'd (LBA)	54	OnTrackDM6	a6	OpenBSD	ef	EFI (FAT-12/16/
10	OPUS	55	EZ-Drive	a7	NeXTSTEP	f0	Linux/PA-RISC b
11	Hidden FAT12	56	Golden Bow	a8	Darwin UFS	f1	SpeedStor
12	Compaq diagnost	5c	Priam Edisk	a9	NetBSD	f4	SpeedStor
14	Hidden FAT16	61	SpeedStor	ab	Darwin boot	f2	DOS secondary
16	Hidden FAT16	63	GNU HURD or Sys	af	HFS / HFS+	fb	VMware VMFS
17	Hidden HPFS/NTF	64	Novell Netware	b7	BSDI fs	fc	VMware VMKCORE
18	AST SmartSleep	65	Novell Netware	b8	BSDI swap	fd	Linux raid auto
1b	Hidden W95 FAT3	70	DiskSecure Mult	bb	Boot Wizard hid	fe	LANstep
1c	Hidden W95 FAT3	75	PC/IX	be	Solaris boot	ff	BBT
1e	Hidden W95 FAT1	80	Old Minix				

Hex code (type L to list all codes): 8e

Changed type of partition 'Linux' to 'Linux LVM'

Command (m for help): w

The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.

The kernel still uses the old table. The new table will be used at

the next reboot or after you run `partprobe(8)` or `kpartx(8)`

Syncing disks.

This error message is expected. You need to reload the Linux machine to have new changes take effect.

```
[root@pcprdusecondary ~]# df -h
```

```
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/rhel-root  4.0G  946M  3.1G  24% /
devtmpfs        3.9G    0  3.9G   0% /dev
tmpfs           3.9G    0  3.9G   0% /dev/shm
tmpfs           3.9G  8.6M  3.9G   1% /run
tmpfs           3.9G    0  3.9G   0% /sys/fs/cgroup
/dev/sda1       1014M  143M  872M  15% /boot
/dev/mapper/rhel-home  2.0G   33M  2.0G   2% /home
tmpfs           781M    0  781M   0% /run/user/0
```

```
[root@pcprdusecondary ~]# fdisk -l
```

Disk /dev/sda: 107.4 GB, 107374182400 bytes, 209715200 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk label type: dos

Disk identifier: 0x00025a26

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	2048	2099199	1048576	83	Linux
/dev/sda2		2099200	31211519	14556160	8e	Linux LVM
/dev/sda3		31211520	209715199	89251840	8e	Linux LVM

Disk /dev/mapper/rhel-root: 4294 MB, 4294967296 bytes, 8388608 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk /dev/mapper/rhel-swap: 8455 MB, 8455716864 bytes, 16515072 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk /dev/mapper/rhel-home: 2147 MB, 2147483648 bytes, 4194304 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

- Create physical volume for sda3.

```
[root@pcprdusecondary ~]# pvcreate /dev/sda3
[root@pcprdusecondary ~]# pvcreate /dev/sda3
Physical volume "/dev/sda3" successfully created.
[root@pcprdusecondary ~]#
```

- pvscan – scan and list physical volume group.
- vgscan – scan and list logical volume group.
- lvscan – scan and list logical volumes created under volume group

This Linux LVM creation is the prerequisite for RDU server installation.

- 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 these specifications:

1. <logical volume for Prime Cable Provisioning install directory> - Mounted on /bprHome directory. For example, LVBPRHOME.

2. <logical volume for Prime Cable Provisioning data directory> - Mounted on /bprData directory. For example, LVBPRDATA

3. <logical volume for Prime Cable Provisioning log directory > - Mounted on /bprLog directory. For example, LVBPRDBLOG

- Create volume group and logical volumes as per requirement and mount on directory /bprData, bprHome and /bprLog directories.

For eg: This procedure is to create logical volumes for BPRHOME with 3 GB disk space, BPRDATA with 15 GB disk space and BPRDBLOG with 5 GB of disk space allocation. You need to choose disk space to extend based on allocation.

- Create Volume group.


```

[root@pcprdusecondary ~]# pvscan
  PU /dev/sda2   UG rhel          lvm2 [<13.88 GiB / 4.00 MiB free]
  PU /dev/sda3           lvm2 [<85.12 GiB]
  Total: 2 [<99.00 GiB] / in use: 1 [<13.88 GiB] / in no UG: 1 [<85.
[root@pcprdusecondary ~]# vgcreate rdusesecondary /dev/sda3
  Volume group "rdusesecondary" successfully created
[root@pcprdusecondary ~]#
[root@pcprdusecondary ~]#
[root@pcprdusecondary ~]# vgs
  Reading volume groups from cache.
  Found volume group "rhel" using metadata type lvm2
  Found volume group "rdusesecondary" using metadata type lvm2
[root@pcprdusecondary ~]# pvscan
  PU /dev/sda2   UG rhel          lvm2 [<13.88 GiB / 4.00 MiB free]
  PU /dev/sda3   UG rdusesecondary lvm2 [85.11 GiB / 85.11 GiB free]
  Total: 2 [98.99 GiB] / in use: 2 [98.99 GiB] / in no UG: 0 [0 ]
[root@pcprdusecondary ~]#

```

vgcreate <vg_name> <pvname>

```

[root@pcprdusecondary ~]# vgcreate rdusesecondary /dev/sda3

```

- Create logical volumes:

lvcreate -L <valuein GB> -n <logicalvolumename> <volumegroupName>

```

[root@pcprdusecondary ~]# lvcreate -L +3GB -n LVBPRHOME rdusesecondary
[root@pcprdusecondary ~]# lvcreate -L +15GB -n LVBPRDATA rdusesecondary
[root@pcprdusecondary ~]# lvcreate -L +5GB -n LVBPRDBLOG rdusesecondary

```

```

[root@pcprdusecondary ~]#
[root@pcprdusecondary ~]# lvcreate -L +3GB -n LVBPRHOME rdusesecondary
  Logical volume "LVBPRHOME" created.
[root@pcprdusecondary ~]# lvcreate -L +15GB -n LVBPRDATA rdusesecondary
  Logical volume "LVBPRDATA" created.
[root@pcprdusecondary ~]# lvcreate -L +5GB -n LVBPRDBLOG rdusesecondary
  Logical volume "LVBPRDBLOG" created.
[root@pcprdusecondary ~]#
[root@pcprdusecondary ~]# LUSCAN
-bash: LUSCAN: command not found
[root@pcprdusecondary ~]# lvs
  ACTIVE          /dev/rhel/root' [4.00 GiB] inherit
  ACTIVE          /dev/rhel/home' [2.00 GiB] inherit
  ACTIVE          /dev/rhel/swap' [<7.00 GiB] inherit
  ACTIVE          /dev/rdusesecondary/LVBPRHOME' [3.00 GiB] inherit
  ACTIVE          /dev/rdusesecondary/LVBPRDATA' [15.00 GiB] inherit
  ACTIVE          /dev/rdusesecondary/LVBPRDBLOG' [5.00 GiB] inherit
[root@pcprdusecondary ~]# vgs
  Reading volume groups from cache.
  Found volume group "rhel" using metadata type lvm2
  Found volume group "rdusesecondary" using metadata type lvm2
[root@pcprdusecondary ~]# pvscan
  PU /dev/sda2   UG rhel          lvm2 [<13.88 GiB / 4.00 MiB free]
  PU /dev/sda3   UG rdusesecondary lvm2 [85.11 GiB / 62.11 GiB free]
  Total: 2 [98.99 GiB] / in use: 2 [98.99 GiB] / in no UG: 0 [0 ]

```

bprHome - installation application path (default directory - /opt/CSCObac)

bprData - installation data path.(Default directory - /var/CSCObac)

bprLog - installation log path. (Default directory - /var/CSCObac)

- Create XFS File system on lvm partition.

mkfs.xfs /dev/<volumegroupName>/<logicalvolume>

```
[root@pcprdusecondary ~]# mkfs.xfs /dev/rdusecondary/LVBPRHOME
[root@pcprdusecondary ~]# mkfs.xfs /dev/rdusecondary/LVBPRDATA
[root@pcprdusecondary ~]# mkfs.xfs /dev/rdusecondary/LVBPRDBLOG
```

```
--- Logical volume ---
LU Path                /dev/rdusecondary/LVBPRDATA
LU Name                LVBPRDATA
UG Name                rdusecondary
LU UUID                d10WKKX-lzuX-NzsY-zSAH-8s8T-qzq6-JM7bn
LU Write Access        read/write
LU Creation host, time pcprdusecondary.cisco.com, 2020-12-02 06:32:25 +053
LU Status               available
# open                 0
LU Size                15.00 GiB
Current LE             3840
Segments               1
Allocation              inherit
Read ahead sectors     auto
- currently set to     8192
Block device           253:4

--- Logical volume ---
LU Path                /dev/rdusecondary/LVBPRDBLOG
LU Name                LVBPRDBLOG
UG Name                rdusecondary
LU UUID                Hd1xm8-jSsf-m6Ax-tUdW-FWz-6k3G-x6zChT
LU Write Access        read/write
LU Creation host, time pcprdusecondary.cisco.com, 2020-12-02 06:34:05 +053
LU Status               available
# open                 0
LU Size                5.00 GiB
Current LE             1200
Segments               1
Allocation              inherit
Read ahead sectors     auto
- currently set to     8192
Block device           253:5
```

- Create directory – bprHome, bprData, bprLog and mount logical volumes on these directories.

```
[root@pcprdusecondary ~]# mkdir bprHome
[root@pcprdusecondary ~]# mkdir bprData
[root@pcprdusecondary ~]# mkdir bprLog
```

- Mount logical volume created on these directories.

```
[root@pcprdusecondary ~]# mount /dev/RDUPRIMARY/LVBPRHOME /bprHome/
[root@pcprdusecondary ~]# mount /dev/RDUPRIMARY/LVBPRDATA /bprData/
[root@pcprdusecondary ~]# mount /dev/RDUPRIMARY/LVBPRDBLOG /bprLog
```

```
[root@pcprdusecondary ~]# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/mapper/rhel-root	4.8G	947M	3.1G	24%	/
devtmpfs	3.9G	0	3.9G	0%	/dev
tmpfs	3.9G	0	3.9G	0%	/dev/shm
tmpfs	3.9G	8.6M	3.9G	1%	/run
tmpfs	3.9G	0	3.9G	0%	/sys/fs/cgroup
/dev/mapper/rhel-home	2.8G	33M	2.8G	2%	/home
/dev/sda1	1014M	143M	872M	15%	/boot
tmpfs	781M	0	781M	0%	/run/user/0
/dev/mapper/rdusecondary-LVBPRHOME	3.8G	33M	3.8G	2%	/bprHome
/dev/mapper/rdusecondary-LVBPRDATA	15G	33M	15G	1%	/bprData
/dev/mapper/rdusecondary-LVBPRDBLOG	5.8G	33M	5.8G	1%	/bprLog

- These commands can be used to enlist and verify new partition status, new physical and logical volume status, file system type, allocation blocks.

```
[root@pcprdusecondary ~]# fdisk -l
[root@pcprdusecondary ~]# pvdisplay
[root@pcprdusecondary ~]# vgdisplay
[root@pcprdusecondary ~]# lvdisplay
```

Note:

- The fstab entries for the logical volumes need not be added. The Corosync cluster will take care of mounting the volumes. In the past, a few customers have faced issues due to these entries. During reboot of the system sometimes due to timing issue both primary and secondary would try to mount the volumes.
- Volume group name and Logical volumes (LVBPRHOME, LVBPRDATA, LVBPRDBLOG) must be the same across both servers. They should share the same disk space on both servers.
- DRBD block device filesystem sync operates only disk size is same on both servers.
- CentOS Linux version must be 7.4 and kernel must be 3.10.0-693.11.6.el7.x86_64.
- Ensure both servers use the same interface for the public IP address where VIP is advertised – ens192.

2. Prepare Linux 7.4 server for RDU HA Deployment on Both Servers

- [RDU HA Installation Mode](#)
- [Common Initial Steps for Configuring RDU HA Nodes](#)
- [RDU HA Setup in Primary-Secondary Mode](#)
- [Preparing RDU Nodes for HA Setup in Primary-Secondary Mode](#)

3. Install RDU server in Geo-Redundancy Mode

- [Setting Up RDU Two Node Failover Pair](#)
- [RDU HA Setup in Primary-Only and Secondary-Only Modes](#)
- [Recovering an Impacted RDU Node Using Recovery Mode](#)

Refer to the quick start guide for more information:

https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/cable_provisioning/6-1-3/quick/start/guide/CiscoPrimeCableProvisioning-6_1_3-QuickStartGuide/CiscoPrimeCableProvisioning-6_1_3-QuickStartGuide_chapter_0101.html#task_1DBF800D2FF84D73BD972A0C6C7B92E6

4. Layer-3 Routing Prerequisite for Geo-Redundancy Deployment

RDU Geo-Redundancy

RDU Geo Redundancy is an enhanced feature of RDU HA supported on RHEL 7.4 or CentOS 7.4 (both 64bit), wherein the RDU primary and secondary node can be in different geographical location or both the nodes can be in a 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.
- In Geo redundancy mode, the VIP will be monitored using the resource agent (res_VIPArp).

PCP Geo-Redundancy Requirements

Route injection for Virtual IP (VIP) 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.

How to redistribute and advertise the RIP2 route to OSPF IGRP. The same redistribution can be used for other protocol like EIGRP/IBGP.

For PCP Geo-redundancy solution the CIDR value of VIP should be 32.

- If VIP advertisement through quagga is enabled then enter the interface through which you want to advertise the VIP, by default it is eth0, make sure this interface name is the same on both primary and secondary servers, also make sure this interface is connected to the ingress router where route injection is done.
- If VIP advertisement through quagga is disabled then enter the CIDR value for VIP
- **/etc/quagga/ripd.conf.** – path where RIP2 conf is added in Geo mode.
<https://www.nongnu.org/quagga/docs/quagga.html#RIP>
- RIP adjacency has to be injected in the neighbouring Router connected to both the primary and secondary server. Sample config like this:

```
Router#show run | sec rip
router rip
version 2
network 10.0.0.0
no auto-summary
Router#_
```

- Adjacency configuration for neighbour peer. This need to implement in both router. VIP and public IP network must be added to advertise interface.
- Route to VIP address.
- Advertise this RIP network via ospf/eigrp/static based on route enabled to advertise to the outside world.

Example: Here OSPF is the dynamic protocol
router ospf <processed>
redistribute rip metric-type 1 subnets. For RIP2, it uses metric as hop count.
Example: Here ISIS is the dynamic protocol
router isis
redistribute rip metric

Post-Check HA

- Check the RDU HA cluster status with the use of the command: **/bprHome/CSCObac/agent/HA/bin/monitor_ha_cluster.sh**.
- Ensure RDU HA works without any issue with Geo-Redundancy mode. Wait for the primary and secondary DRBD disks to sync and show up-to-date status (cat /proc/drbd).