



Red Hat Enterprise Linux

5

Configuration Example - Fence Devices

Configuring Fence Devices in a Red Hat Cluster
Edition 2

Landmann

Red Hat Enterprise Linux 5 Configuration Example - Fence Devices

Configuring Fence Devices in a Red Hat Cluster Edition 2

Landmann
rlandmann@redhat.com

Legal Notice

Copyright © 2010 Red Hat Inc..

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This book describes a procedure for configuring fence devices in a Red Hat Cluster using Conga.

Table of Contents

Chapter 1. Introduction	3
1.1. About This Guide	3
1.2. Audience	3
1.3. Software Versions	3
1.4. Related Documentation	3
Chapter 2. Configuring Fence Devices in a Red Hat Cluster	5
Chapter 3. Configuring an APC Switch as a Fence Device	6
3.1. APC Fence Device Prerequisite Configuration	6
3.2. APC Fence Device Components to Configure	7
3.3. APC Fence Device Configuration Procedure	8
3.4. Cluster Configuration File with APC Fence Device	11
3.5. Testing the APC Fence Device Configuration	12
Chapter 4. Configuring IPMI Management Boards as Fencing Devices	13
4.1. IPMI Fence Device Prerequisite Configuration	13
4.2. IPMI Fence Device Components to Configure	14
4.3. IPMI Fence Device Configuration Procedure	15
4.4. Cluster Configuration File with IPMI Fence Device	17
4.5. Testing the IPMI Fence Device Configuration	19
Chapter 5. Configuring HP iLO Management Boards as Fencing Devices	20
5.1. HP iLO Fence Device Prerequisite Configuration	20
5.2. HP iLO Fence Device Components to Configure	21
5.3. HP iLO Fence Device Configuration Procedure	22
5.4. Cluster Configuration File with HP iLO Fence Device	24
5.5. Testing the HP iLO Fence Device Configuration	26
Chapter 6. Configuring Fencing with Dual Power Supplies	27
6.1. Dual Power Fencing Prerequisite Configuration	27
6.2. Fence Device Components to Configure	28
6.3. Dual Power Fencing Configuration Procedure	29
6.4. Cluster Configuration File with Dual Power Supply Fencing	34
6.5. Testing the Dual Power Fence Device Configuration	36
Chapter 7. Configuring a Backup Fencing Method	37
7.1. Backup Fencing Prerequisite Configuration	37
7.2. Fence Device Components to Configure	39
7.3. Backup Fencing Configuration Procedure	41
7.4. Cluster Configuration File for Backup Fence Method	48
7.5. Testing the Backup Fence Device Configuration	50
Chapter 8. Configuring Fencing using SCSI Persistent Reservations	51
8.1. Technical Overview of SCSI Persistent Reservations	51
8.2. SCSI Fencing Requirements and Limitations	51
8.3. SCSI Fencing Example Configuration	52
8.4. SCSI Fencing Prerequisite Configuration	52
8.5. SCSI Fence Device Components to Configure	53
8.6. SCSI Fence Device Configuration Procedure	54
8.7. Cluster Configuration File with SCSI Fence Device	55
8.8. Testing the Configuration	56
Chapter 9. Troubleshooting	58

Chapter 10. The GFS Withdraw Function	61
Appendix A. Revision History	62
Index	62

Chapter 1. Introduction

1.1. About This Guide

This book describes procedures for configuring fence devices in a Red Hat Cluster using Conga.

1.2. Audience

This book is intended to be used by system administrators managing systems running the Linux operating system. It requires familiarity with Red Hat Enterprise Linux 5 and Red Hat Cluster Suite.

1.3. Software Versions

Table 1.1. Software Versions

Software	Description
RHEL5	refers to Red Hat Enterprise Linux 5 and higher
GFS	refers to GFS for Red Hat Enterprise Linux 5 and higher

1.4. Related Documentation

For more information about using Red Hat Enterprise Linux, refer to the following resources:

- ✦ *Red Hat Enterprise Linux Installation Guide* — Provides information regarding installation of Red Hat Enterprise Linux 5.
- ✦ *Red Hat Enterprise Linux Deployment Guide* — Provides information regarding the deployment, configuration and administration of Red Hat Enterprise Linux 5.

For more information about Red Hat Cluster Suite for Red Hat Enterprise Linux 5, refer to the following resources:

- ✦ *Red Hat Cluster Suite Overview* — Provides a high level overview of the Red Hat Cluster Suite.
- ✦ *Configuring and Managing a Red Hat Cluster* — Provides information about installing, configuring and managing Red Hat Cluster components.
- ✦ *Logical Volume Manager Administration* — Provides a description of the Logical Volume Manager (LVM), including information on running LVM in a clustered environment.
- ✦ *Global File System: Configuration and Administration* — Provides information about installing, configuring, and maintaining Red Hat GFS (Red Hat Global File System).
- ✦ *Global File System 2: Configuration and Administration* — Provides information about installing, configuring, and maintaining Red Hat GFS2 (Red Hat Global File System 2).
- ✦ *Using Device-Mapper Multipath* — Provides information about using the Device-Mapper Multipath feature of Red Hat Enterprise Linux 5.
- ✦ *Using GNBD with Global File System* — Provides an overview on using Global Network Block Device (GNBD) with Red Hat GFS.

- ✦ *Linux Virtual Server Administration* — Provides information on configuring high-performance systems and services with the Linux Virtual Server (LVS).
- ✦ *Red Hat Cluster Suite Release Notes* — Provides information about the current release of Red Hat Cluster Suite.

Red Hat Cluster Suite documentation and other Red Hat documents are available in HTML, PDF, and RPM versions on the Red Hat Enterprise Linux Documentation CD and online at <http://www.redhat.com/docs/>.

Chapter 2. Configuring Fence Devices in a Red Hat Cluster

This document provides configuration examples that show the steps needed to configure fence devices in a Red Hat cluster using the Conga configuration tool. For general information about fencing and fence device configuration, see *Configuring and Managing a Red Hat Cluster*.

This remainder of this document is organized as follows:

- [Chapter 3, Configuring an APC Switch as a Fence Device](#) describes the procedure for configuring an APC switch as a fence device in a Red Hat cluster.
- [Chapter 4, Configuring IPMI Management Boards as Fencing Devices](#) describes the procedure for configuring IPMI management boards as fence devices in a Red Hat cluster.
- [Chapter 5, Configuring HP iLO Management Boards as Fencing Devices](#) describes the procedure for configuring HP iLO management boards as fence devices in a Red Hat cluster.
- [Chapter 6, Configuring Fencing with Dual Power Supplies](#) describes the procedure for configuring two APC switches using separate power supplies to fence each cluster node in a Red Hat cluster.
- [Chapter 7, Configuring a Backup Fencing Method](#) describes the procedure for configuring two APC switches using separate power supplies as a main fencing method and a separate IPMI management board as a backup fencing method to fence each cluster node in a Red Hat cluster.
- [Chapter 8, Configuring Fencing using SCSI Persistent Reservations](#) describes the procedure for configuring fencing on a system using SCSI persistent reservations in a Red Hat cluster.
- [Chapter 9, Troubleshooting](#) provides some guidelines to follow when your configuration does not behave as expected.
- [Chapter 10, The GFS Withdraw Function](#) summarizes some general concerns to consider when configuring fence devices in a Red Hat cluster.

Chapter 3. Configuring an APC Switch as a Fence Device

This chapter provides the procedures for configuring an APC switch as a fence device in a Red Hat cluster using the Conga configuration tool.

[Figure 3.1, “Using an APC Switch as a Fence Device”](#) shows the configuration this procedure yields. In this configuration a three node cluster uses an APC switch as the fencing device. Each node in the cluster is connected to a port in the APC switch.

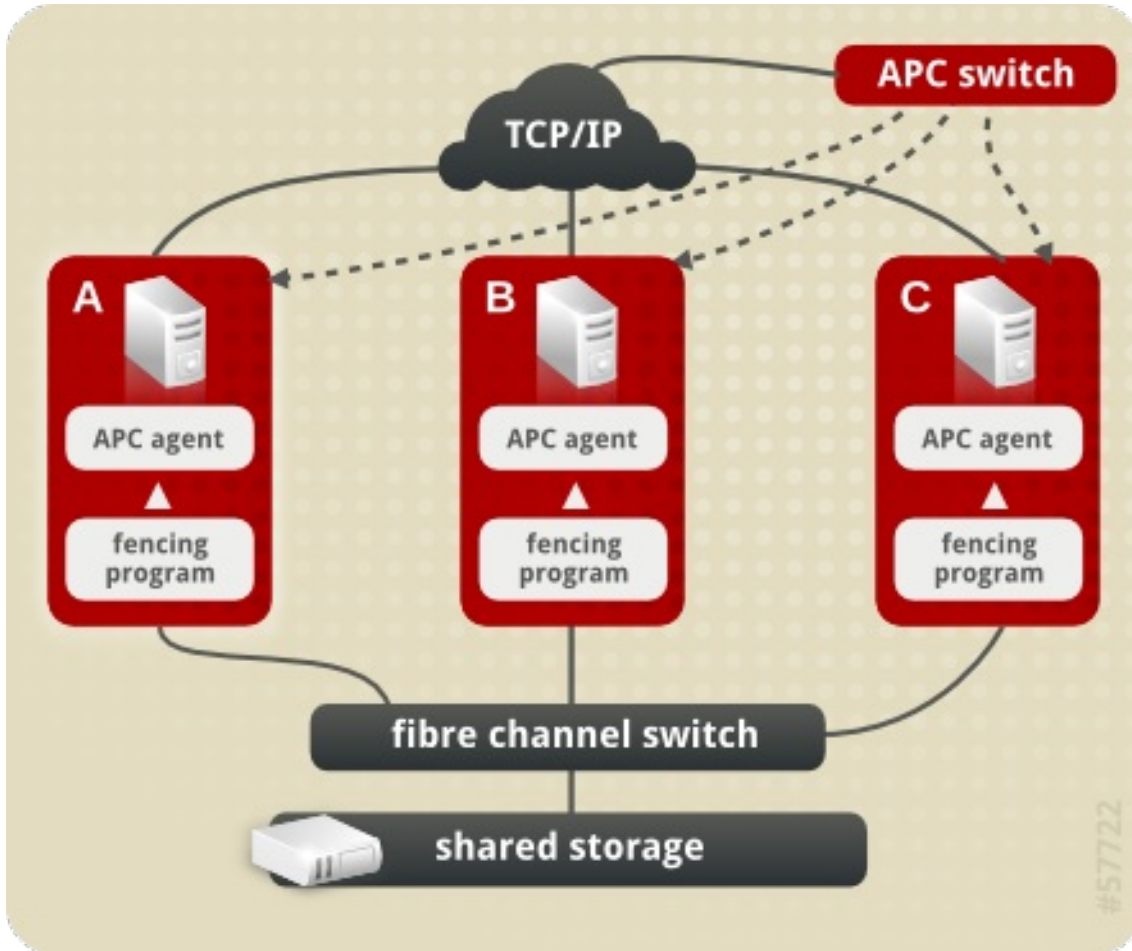


Figure 3.1. Using an APC Switch as a Fence Device

3.1. APC Fence Device Prerequisite Configuration

[Table 3.1, “Configuration Prerequisites”](#) summarizes the prerequisite components that have been set up before this procedure begins.

Table 3.1. Configuration Prerequisites

Component	Name	Comment
cluster	apclust	three-node cluster
cluster node	clusternode1.example.com	node in cluster apclust configured with APC switch to administer power supply
cluster node	clusternode2.example.com	node in cluster apclust configured with APC switch to administer power supply
cluster node	clusternode3.example.com	node in cluster apclust configured with APC switch to administer power supply

Component	Name	Comment
IP address	10.15.86.96	IP address for the APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
login	apclogin	login value for the APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
password	apcpword	password for the APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
port	1	port number on APC switch that clusternode1.example.com connects to
port	2	port number on APC switch that clusternode2.example.com connects to
port	3	port number on APC switch that clusternode3.example.com connects to

3.2. APC Fence Device Components to Configure

This procedure configures an APC switch as a fence device that will be used for each node in cluster **apcclust**. Then the procedure configures that switch as the fencing device for **clusternode1.example.com**, **clusternode2.example.com**, and **clusternode1.example.com**.

[Table 3.2, “Fence Device Components to Configure for APC Fence Device”](#) summarizes the components of the APC fence device that this procedure configures for each of the cluster nodes in **clusternode1.example.com**.

Table 3.2. Fence Device Components to Configure for APC Fence Device

Fence Device Component	Value	Description
Fencing Type	APC Power Switch	type of fencing device to configure
Name	apcfence	name of the APC fencing device
IP address	10.15.86.96	IP address of the APC switch to configure as a fence device for node1.example.com, node2.example.com, and node3.example.com
login	apclogin	login value for the APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
password	apcpword	password for the APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com

[Table 3.3, “Fence Agent Components to Specify for Each Node in apcclust”](#) summarizes the components of the APC fence device that you must specify for the cluster nodes in **apcclust**.

Table 3.3. Fence Agent Components to Specify for Each Node in apcclust

Fence Agent Component	Value	Description
fence device	apcfence	name of the APC fence device you defined as a shared device
port	1	port number on the APC switch for node1.example.com
port	2	port number on the APC switch for node2.example.com
port	3	port number on the APC switch for node3.example.com

The remainder of the fence device components that you configure for each node appear automatically when you specify that you will be configuring the **apcfence** fence device that you previously defined as a shared fence device.

3.3. APC Fence Device Configuration Procedure

This section provides the procedure for adding an APC fence device to each node of cluster **apcclust**. This example uses the same APC switch for each cluster node. The APC fence device will first be configured as a shared fence device. After configuring the APC switch as a shared fence device, the device will be added as a fence device for each node in the cluster.

To configure an APC switch as a shared fence device using **Conga**, perform the following procedure:

1. As an administrator of **luci** Select the **cluster** tab. This displays the **Choose a cluster to administer** screen.
2. From the **Choose a cluster to administer** screen, you should see the previously configured cluster **apcclust** displayed, along with the nodes that make up the cluster. Click on **apcclust** to select the cluster.
3. At the detailed menu for the cluster **apcclust** (below the **clusters** menu on the left side of the screen), click **Shared Fence Devices**. Clicking **Shared Fence Devices** causes the display of any shared fence devices previously configured for a cluster and causes the display of menu items for fence device configuration: **Add a Fence Device** and **Configure a Fence Device**.
4. Click **Add a Fence Device**. Clicking **Add a Fence Device** causes the **Add a Sharable Fence Device** page to be displayed.
5. At the **Add a Sharable Fence Device** page, click the drop-down box under **Fencing Type** and select **APC Power Switch**. This causes Conga to display the components of an APC Power Switch fencing type, as shown in [Figure 3.2, "Adding a Sharable Fence Device"](#).

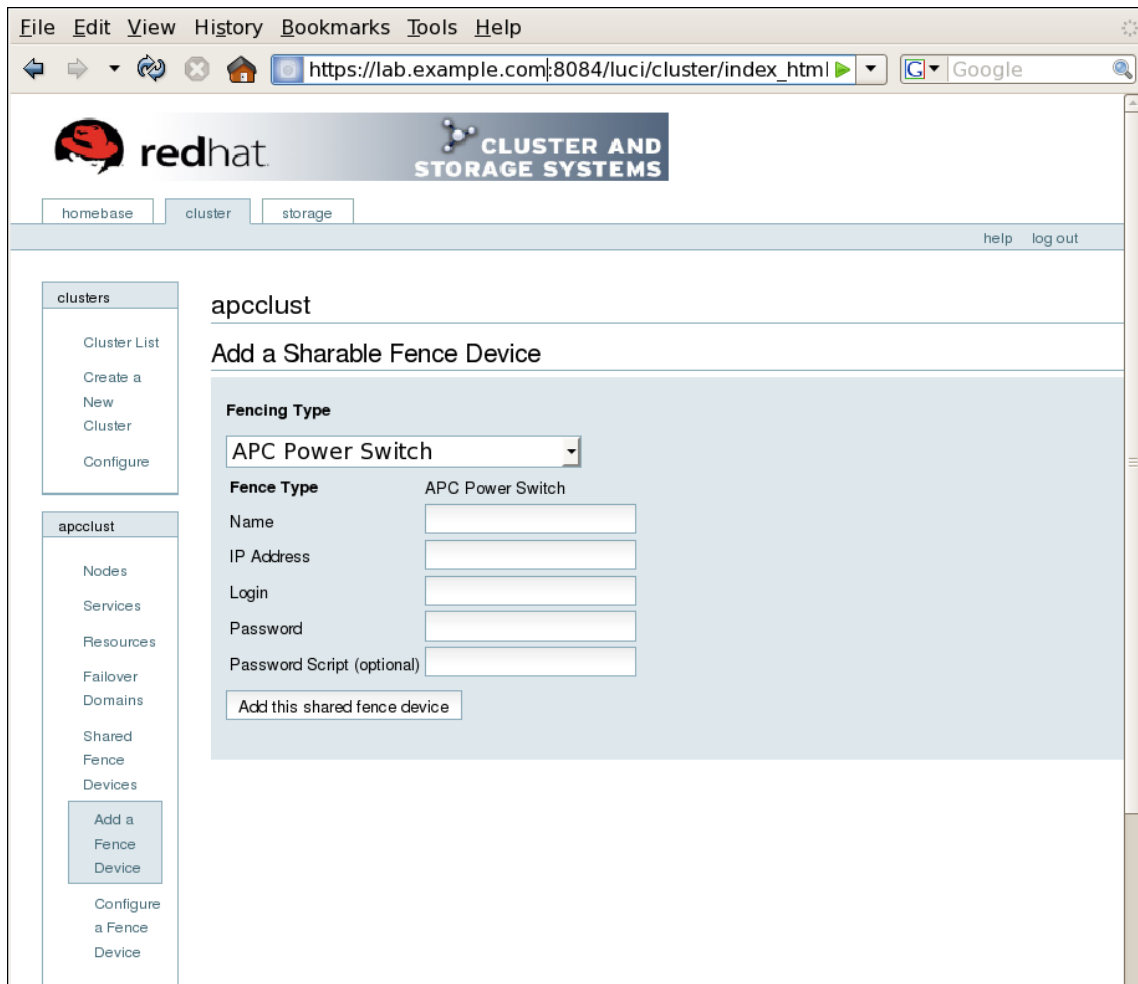


Figure 3.2. Adding a Sharable Fence Device

6. For **Name**, enter **apcfence**.
7. For **IP Address**, enter **10.15.86.96**.
8. For **Login**, enter **apclogin**.
9. For **Password**, enter **apcpword**.
10. For **Password Script**, leave blank.
11. Click **Add this shared fence device**.

Clicking **Add this shared fence device** causes a progress page to be displayed temporarily. After the fence device has been added, the detailed cluster properties menu is updated with the fence device under **Configure a Fence Device**.

After configuring the APC switch as a shared fence device, use the following procedure to configure the APC switch as the fence device for node **clusternode1.example.com**

1. At the detailed menu for the cluster **apcclust** (below the **clusters** menu), click **Nodes**. Clicking **Nodes** causes the display of the status of each node in **apcclust**.
2. At the bottom of the display for node **clusternode1.example.com**, click **Manage Fencing for this Node**. This displays the configuration screen for node **clusternode1.example.com**.

- At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
- From the dropdown menu, the **apcfence** fence device you have already created should display as one of the menu options under **Use an Existing Fence Device**. Select **apcfence (APC Power Device)**. This causes a fence device configuration menu to display with the **Name, IP Address, Login, Password, and Password Script** values already configured, as defined when you configured **apcfence** as a shared fence device. This is shown in [Figure 3.3, “Adding an Existing Fence Device to a Node”](#).

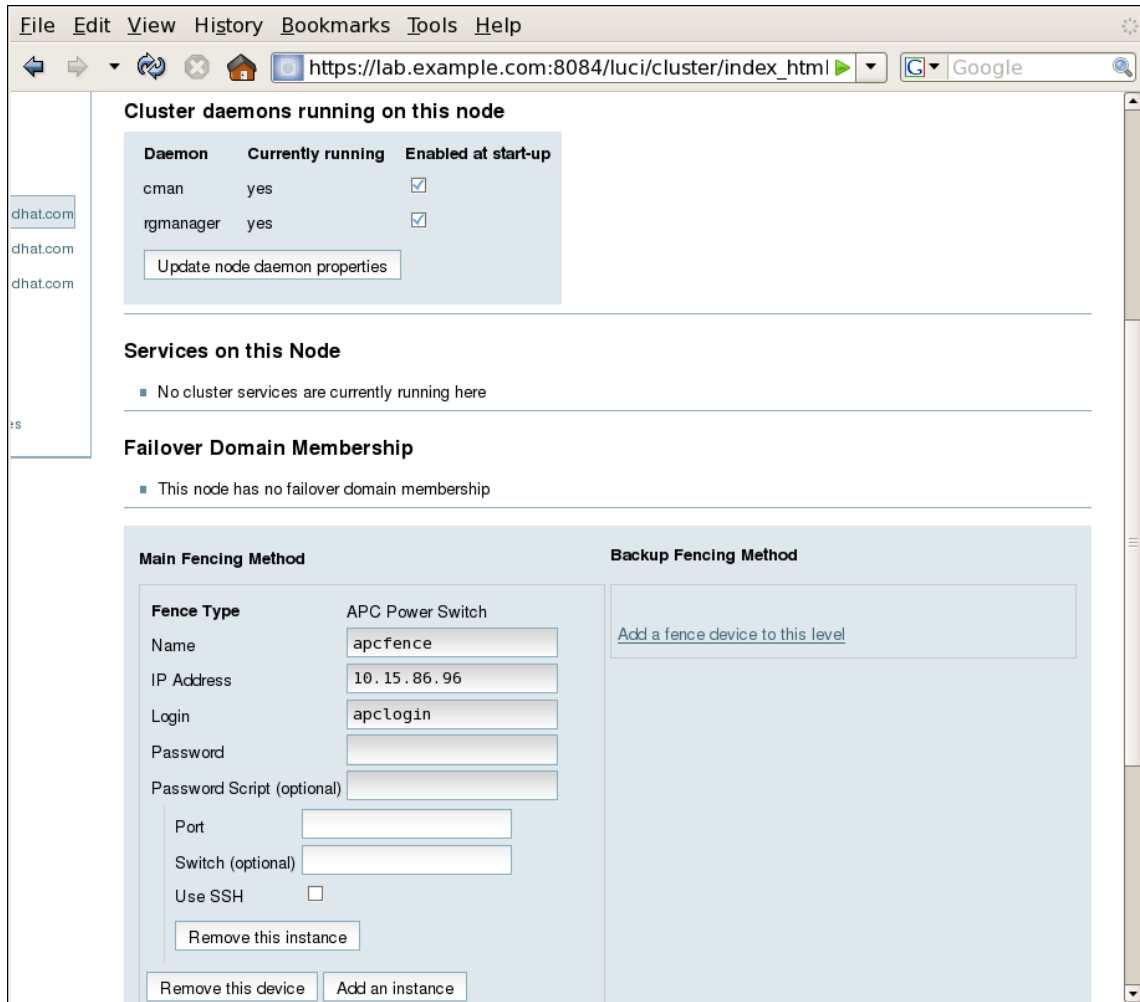


Figure 3.3. Adding an Existing Fence Device to a Node

- For **Port**, enter **1**. Do not enter any value for **Switch**.
- Click **Update main fence properties**. This causes a confirmation screen to be displayed.
- On the confirmation screen, Click **OK**. A progress page is displayed after which the display returns to the status page for **clusternode1.example.com** in cluster **apcclust**.

After configuring **apcfence** as the fencing device for **clusternode1.example.com**, use the same procedure to configure **apcfence** as the fencing device for **clusternode2.example.com**, specifying Port 2 for **clusternode2.example.com**, as in the following procedure:

- On the status page for **clusternode1.example.com** in cluster **apcclust**, the other nodes in **apcclust** are displayed below the **Configure** menu item below the **Nodes** menu item on the left side of the screen. Click **clusternode2.example.com** to display the status screen for **clusternode2.example.com**.

2. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
3. As for **clusternode1.example.com**, the **apcfence** fence device should display as one of the menu options on the dropdown menu, under **Use an Existing Fence Device**. Select **apcfence (APC Power Device)**. This causes a fence device configuration menu to display with the **Name, IP Address, Login, Password, Password Script** values already configured, as defined when you configured **apcfence** as a shared fence device.
4. For **Port**, enter **2**. Do not enter any value for **Switch**.
5. Click **Update main fence properties**.

Similarly, configure **apcfence** as the main fencing method for **clusternode3.example.com**, specifying **3** as the Port number.


3.4. Cluster Configuration File with APC Fence Device

Configuring a cluster with Conga modifies the cluster configuration file. This section shows the cluster configuration file before and after the procedures documented in [Section 3.3, "APC Fence Device Configuration Procedure"](#) were performed.

Before the cluster resources and service were configured, the **cluster.conf** file appeared as follows.

```
<?xml version="1.0"?>
<cluster alias="apcclust" config_version="12" name="apcclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
      <fence/>
    </clusternode>
  </clusternodes>
  <cman/>
  <fencedevices/>
  <rm>
    <failoverdomains/>
    <resources/>
  </rm>
</cluster>
```

After the cluster resources and service were configured, the **cluster.conf** file appears as follows.




```

<?xml version="1.0"?>
<cluster alias="apcclust" config_version="19" name="apcclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence>
        <method name="1">
          <device name="apcfence" port="1"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
      <fence>
        <method name="1">
          <device name="apcfence" port="2"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
      <fence>
        <method name="1">
          <device name="apcfence" port="3"/>
        </method>
      </fence>
    </clusternode>
  </clusternodes>
  <cman/>
  <fencedevices>
    <fencedevice agent="fence_apc" ipaddr="10.15.86.96"
login="apclogin" name="apcfence" passwd="apcpword"/>
  </fencedevices>
  <rm>
    <failoverdomains/>
    <resources/>
  </rm>
</cluster>

```

3.5. Testing the APC Fence Device Configuration

To check whether the configuration you have defined works as expected, you can use the **fence_node** to fence a node manually. The **fence_node** program reads the fencing settings from the **cluster.conf** file for the given node and then runs the configured fencing agent against the node.

To test whether the APC switch has been successfully configured as a fence device for the three nodes in cluster **apcclust**, execute the following commands and check whether the nodes have been fenced.

```

# /sbin/fence_node clusternode1.example.com
# /sbin/fence_node clusternode2.example.com
# /sbin/fence_node clusternode3.example.com

```


Chapter 4. Configuring IPMI Management Boards as Fencing Devices

This chapter provides the procedures for configuring IPMI management boards as fencing devices in a Red Hat cluster using the Conga configuration tool.

[Figure 4.1, “Using IPMI Management Boards as Fence Devices”](#) shows the configuration this procedure yields. In this configuration each node of a three node cluster uses an IPMI management board as its fencing device.



Note

Note that in this configuration each system has redundant power and is hooked into two independent power sources. This ensures that the management board would still function as needed in a cluster even if you lose power from one of the sources.

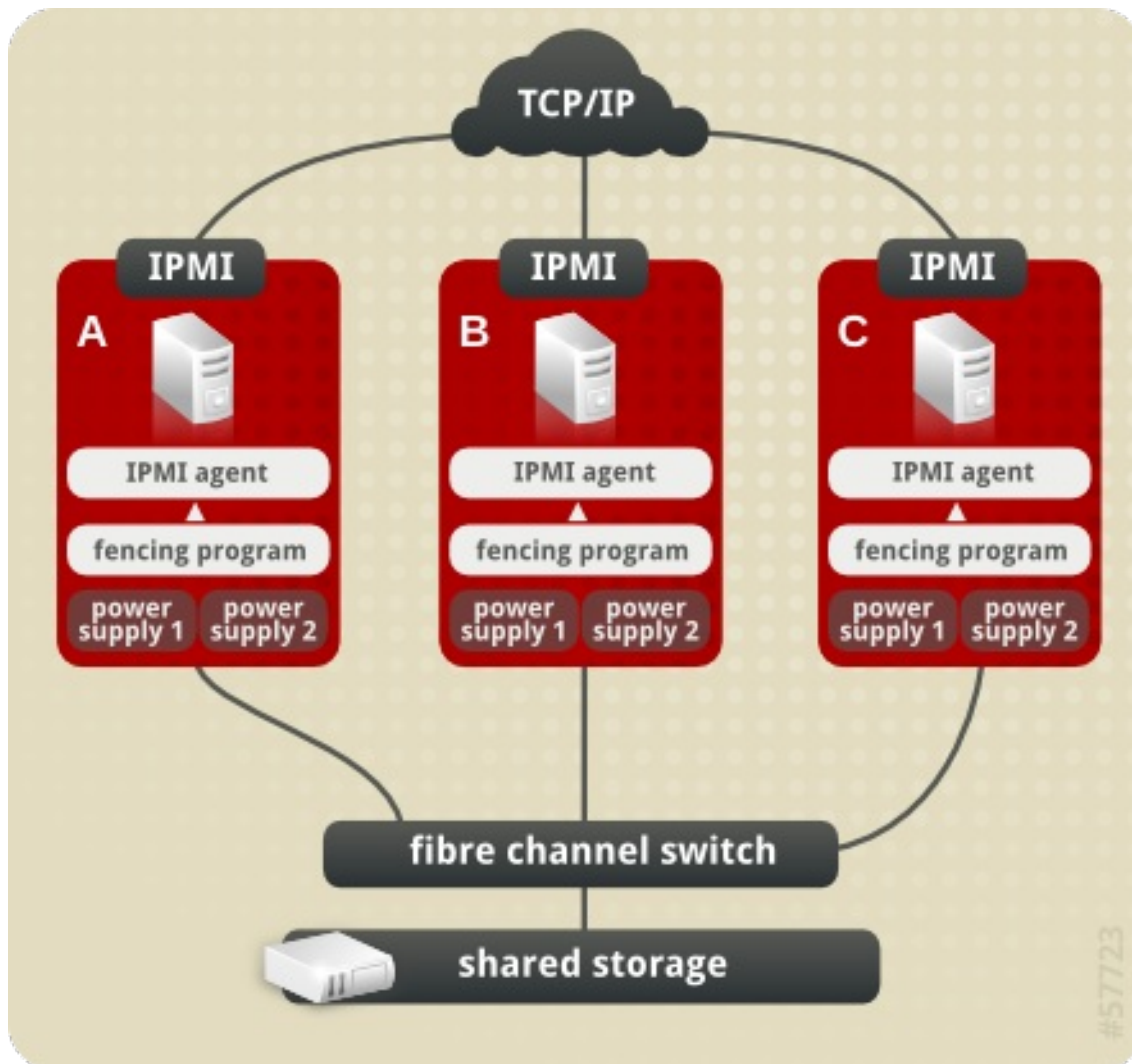


Figure 4.1. Using IPMI Management Boards as Fence Devices

4.1. IPMI Fence Device Prerequisite Configuration

[Table 4.1, “Configuration Prerequisites”](#) summarizes the prerequisite components that have been set up before this procedure begins.

Table 4.1. Configuration Prerequisites

Component	Name	Comment
cluster	ipmiclust	three-node cluster
cluster node	clusternode1.example.com	node in cluster ipmiclust configured with IPMI management board and two power supplies
IP address	10.15.86.96	IP address for IPMI management board for clusternode1.example.com
login	ipmilogin	login name for IPMI management board for clusternode1.example.com
password	ipmipword	password IPMI management board for clusternode1.example.com
cluster node	clusternode2.example.com	node in cluster ipmiclust configured with IPMI management board and two power supplies
IP address	10.15.86.97	IP address for IPMI management board for clusternode2.example.com
login	ipmilogin	login name for IPMI management board for clusternode2.example.com
password	ipmipword	password for IPMI management board for clusternode2.example.com
cluster node	clusternode3.example.com	node in cluster ipmiclust configured with IPMI management board and two power supplies
IP address	10.15.86.98	IP address for IPMI management board for clusternode3.example.com
login	ipmilogin	login name for IPMI management board for clusternode3.example.com
password	ipmipword	password for IPMI management board for clusternode3.example.com

4.2. IPMI Fence Device Components to Configure

This procedure configures the IPMI management board as a fence device for each node in cluster `ipmiclust`.

[Table 4.2, “Fence Agent Components to Configure for clusternode1.example.com”](#) summarizes the components of the IPMI fence device that this procedure configures for cluster node `clusternode1.example.com`.

Table 4.2. Fence Agent Components to Configure for clusternode1.example.com

Fence Agent Component	Value	Description
Name	ipmifence1	name of the IPMI fencing device
IP address	10.15.86.96	IP address of the IPMI management board to configure as a fence device for clusternode1.example.com
IPMI login	ipmilogin	login identity for the IPMI management board for clusternode1.example.com

Fence Agent Component	Value	Description
password	ipmipword	password for the IPMI management board for clusternode1.example.com
authentication type	password	authentication type for the IPMI management board for clusternode1.example.com

[Table 4.3, “Fence Agent Components to Configure for clusternode2.example.com”](#) summarizes the components of the IPMI fence device that this procedure configures for cluster node `clusternode2.example.com`.

Table 4.3. Fence Agent Components to Configure for clusternode2.example.com

Fence Agent Component	Value	Description
Name	ipmifence2	name of the IPMI fencing device
IP address	10.15.86.97	IP address of the IPMI management board to configure as a fence device for clusternode2.example.com
IPMI login	ipmilogin	login identity for the IPMI management board for clusternode2.example.com
password	ipmipword	password for the IPMI management board for clusternode2.example.com
authentication type	password	authentication type for the IPMI management board for clusternode2.example.com

[Table 4.4, “Fence Agent Components to Configure for clusternode3.example.com”](#) summarizes the components of the IPMI fence device that this procedure configures for cluster node `clusternode3.example.com`.

Table 4.4. Fence Agent Components to Configure for clusternode3.example.com

Fence Agent Component	Value	Description
Name	ipmifence3	name of the IPMI fencing device
IP address	10.15.86.98	IP address of the IPMI management board to configure as a fence device for clusternode3.example.com
IPMI login	ipmilogin	login identity for the IPMI management board for clusternode3.example.com
password	ipmipword	password for the IPMI management board for clusternode3.example.com
authentication type	password	authentication type for the IPMI management board for clusternode3.example.com

4.3. IPMI Fence Device Configuration Procedure

This section provides the procedure for adding an IPMI fence device to each node of cluster `ipmiclust`. Each node of `ipmiclust` is managed by its own IPMI management board.

Use the following procedure to configure the IPMI management board as the fence device for node **clusternode1.example.com** using Conga:

1. As an administrator of **luci** Select the **cluster** tab. This displays the **Choose a cluster to administer** screen.
2. From the **Choose a cluster to administer** screen, you should see the previously configured cluster **ipmiclust** displayed, along with the nodes that make up the cluster. Click on **clusternode1.example.com**. This displays the configuration screen for node **clusternode1.example.com**.
3. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
4. From the dropdown menu, under **Create a new Fence Device**, select **IPMI Lan**. This displays a fence device configuration menu, as shown in [Figure 4.2, "Creating an IPMI Fence Device"](#).

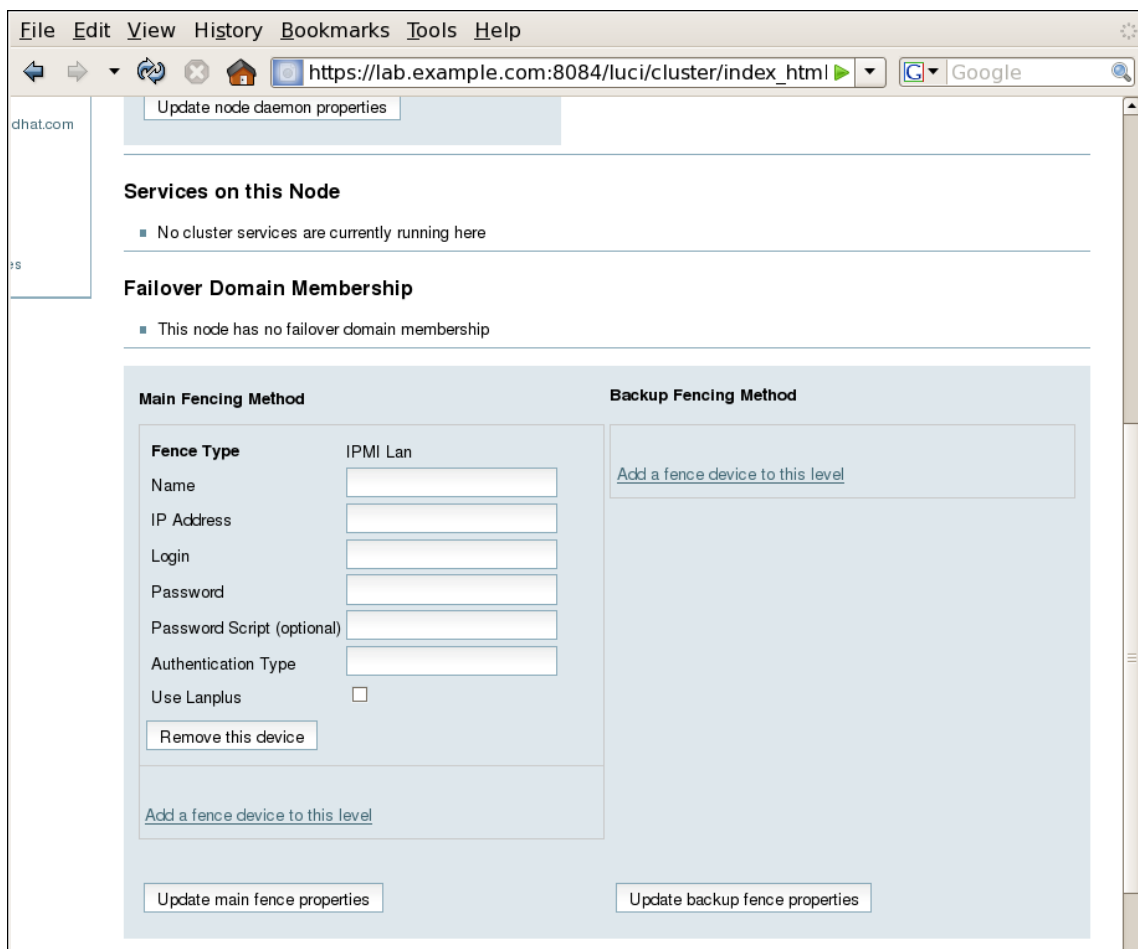


Figure 4.2. Creating an IPMI Fence Device

5. For **Name**, enter **ipmifence1**.
6. For **IP Address**, enter **10.15.86.96**.
7. For **Login**, enter **ipmilogin**.
8. For **Password**, enter **ipmipword**.
9. For **Password Script**, leave the field blank.
10. For **Authentication type**, enter **password**. This field specifies the IPMI authentication type. Possible

values for this field are none, **password**, **md2**, or **md5**.

11. Leave the **Use Lanplus** field blank. You would check this field if your fence device is a Lanplus-capable interface such as iLO2.
12. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
13. On the confirmation screen, click **OK**. After the fence device has been added, a progress page is displayed after which the display returns to the configuration page for **clusternode1.example.com** in cluster **ipmiclust**.

After configuring an IPMI fence device for **clusternode1.example.com**, use the following procedure to configure an IPMI fence device for **clusternode2.example.com**.

1. From the configuration page for **clusternode1.example.com**, a menu appears on the left of the screen for cluster **ipmiclust**. Select the node **clusternode2.example.com**. The configuration page for **clusternode2.example.com** appears, with no fence device configured.
2. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
3. From the dropdown menu, under **Create a new Fence Device**, select **IPMI Lan**. This displays a fence device configuration menu.
4. For **Name**, enter **ipmifence2**.
5. For **IP Address**, enter **10.15.86.97**.
6. For **Login**, enter **ipmilogin**.
7. For **Password**, enter **ipmipword**.
8. For **Password Script**, leave the field blank.
9. For **Authentication type**, enter **password**. This field specifies the IPMI authentication type. Possible values for this field are none, **password**, **md2**, or **md5**.
10. Leave the **Use Lanplus** field blank.
11. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
12. On the confirmation screen, click **OK**. After the fence device has been added, a progress page is displayed after which the display returns to the configuration page for **clusternode1.example.com** in cluster **ipmiclust**.

After configuring **ipmifence2** as the fencing device for **clusternode2.example.com**, select node **clusternode3.example.com** from the menu on the left side of the page and configure an IPMI fence device for that node using the same procedure as you did to configure the fence devices for **clusternode2.example.com** and **clusternode3.example.com**. For **clusternode3.example.com**, use **ipmifence3** as the name of the fencing method and 10.15.86.98 as the IP address. Otherwise, use the same values for the fence device parameters.

4.4. Cluster Configuration File with IPMI Fence Device

Configuring a cluster with Conga modifies the cluster configuration file. This section shows the cluster configuration file before and after the procedures documented in [Section 4.3, "IPMI Fence Device Configuration Procedure"](#) and were performed.

Before the cluster resources and service were configured, the `cluster.conf` file appeared as follows.

```
<?xml version="1.0"?>
<cluster alias="ipmiclust" config_version="12" name="ipmiclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
      <fence>
        <method name="1"/>
      </fence>
    </clusternode>
  </clusternodes>
  <cman/>
  <fencedevices/>
  <rm>
    <failoverdomains/>
    <resources/>
  </rm>
</cluster>
```

After the cluster resources and service were configured, the `cluster.conf` file appears as follows.

```
<?xml version="1.0"?>
<cluster alias="ipmiclust" config_version="27" name="ipmiclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence>
        <method name="1">
          <device name="ipmifence1"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
      <fence>
        <method name="1">
          <device name="ipmifence2"/>
        </method>
      </fence>
    </clusternode>
  </clusternodes>
  <cman/>
  <fencedevices/>
  <rm>
    <failoverdomains/>
    <resources/>
  </rm>
</cluster>
```

```

        <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
            <fence>
                <method name="1">
                    <device name="ipmifence3"/>
                </method>
            </fence>
        </clusternode>
    </clusternodes>
    <cman/>
    <fencedevices>
        <fencedevice agent="fence_ipmilan" ipaddr="10.15.86.96"
login="ipmilogin" name="ipmifence1" passwd="ipmipword" />
        <fencedevice agent="fence_ipmilan" ipaddr="10.15.86.97"
login="ipmilogin" name="ipmifence2" passwd="ipmipword" />
        <fencedevice agent="fence_ipmilan" ipaddr="10.15.86.98"
login="ipmilogin" name="ipmifence3" passwd="ipmipword" />
    </fencedevices>
    <rm>
        <failoverdomains/>
        <resources/>
    </rm>
</cluster>

```

4.5. Testing the IPMI Fence Device Configuration

To check whether the configuration you have defined works as expected, you can use the **fence_node** to fence a node manually. The **fence_node** program reads the fencing settings from the **cluster.conf** file for the given node and then runs the configured fencing agent against the node.

To test whether the IPMI management boards have been successfully configured as fence devices for the three nodes in cluster **ipmiclust**, execute the following commands and check whether the nodes have been fenced.

```

# /sbin/fence_node clusternode1.example.com
# /sbin/fence_node clusternode2.example.com
# /sbin/fence_node clusternode3.example.com

```


Chapter 5. Configuring HP iLO Management Boards as Fencing Devices

This chapter provides the procedures for configuring HP iLO management boards as fencing devices in a Red Hat cluster using the Conga configuration tool.

[Figure 5.1, “Using HP iLO Management Boards as Fence Devices”](#) shows the configuration this procedure yields. In this configuration each node of a three node cluster uses an HP iLO management board as its fencing device.



Note

Note that in this configuration each system has redundant power and is hooked into two independent power sources. This ensures that the management board would still function as needed in a cluster even if you lose power from one of the sources.

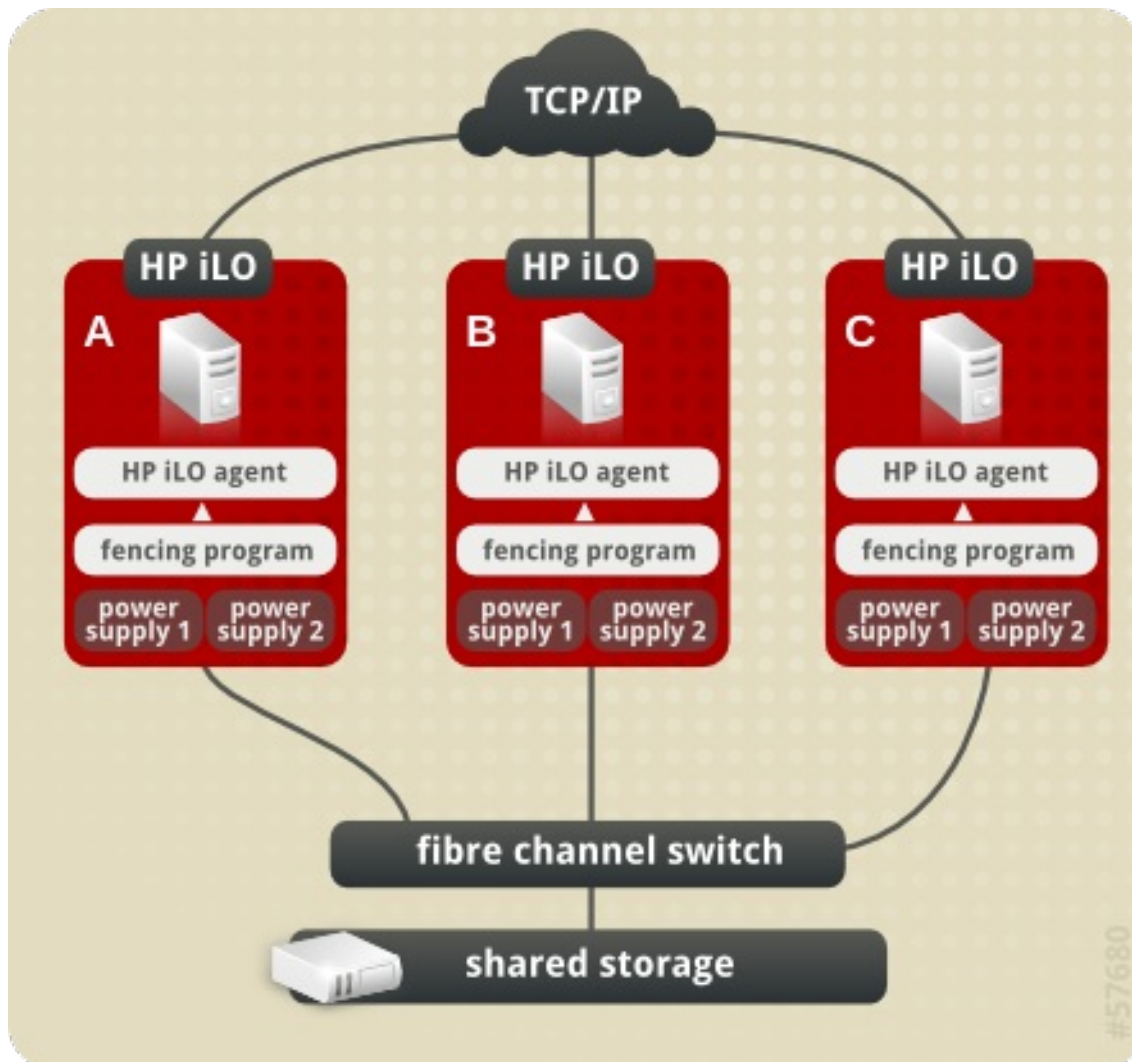


Figure 5.1. Using HP iLO Management Boards as Fence Devices

5.1. HP iLO Fence Device Prerequisite Configuration

[Table 5.1, “Configuration Prerequisites”](#) summarizes the prerequisite components that have been set up before this procedure begins.

Table 5.1. Configuration Prerequisites

Component	Name	Comment
cluster	hpiloclust	three-node cluster
cluster node	clusternode1.example.com	node in cluster hpiloclust configured with HP iLO management board and two power supplies
hostname	hpilohost1	host name for HP iLO management board for clusternode1.example.com
login	hpilologin	login name for HP iLO management board for clusternode1.example.com
password	hpilopword	password HP iLO management board for clusternode1.example.com
cluster node	clusternode2.example.com	node in cluster hpiloclust configured with HP iLO management board and two power supplies
hostname	hpilohost2	hostname for HP iLO management board for clusternode2.example.com
login	hpilologin	login name for HP iLO management board for clusternode2.example.com
password	hpilopword	password for HP iLO management board for clusternode2.example.com
cluster node	clusternode3.example.com	node in cluster hpiloclust configured with HP iLO management board and two power supplies
hostname	hpilohost3	host name for HP iLO management board for clusternode3.example.com
login	hpilologin	login name for HP iLO management board for clusternode3.example.com
password	hpilopword	password for HP iLO management board for clusternode3.example.com

5.2. HP iLO Fence Device Components to Configure

This procedure configures the HP iLO management board as a fence device for each node in cluster **hpiloclust**.

[Table 5.2, “Fence Agent Components to Configure for clusternode1.example.com”](#) summarizes the components of the HP iLO fence device that this procedure configures for cluster node **clusternode1.example.com**.

Table 5.2. Fence Agent Components to Configure for clusternode1.example.com

Fence Agent Component	Value	Description
Name	hpilofence1	name of the HP iLO fencing device
hostname	hpilohost1	host name of the HP iLO management board to configure as a fence device for clusternode1.example.com
HP iLO login	hpilologin	login identity for the HP iLO management board for clusternode1.example.com

Fence Agent Component	Value	Description
password	hpilopword	password for the HP iLO management board for clusternode1.example.com

[Table 5.3, "Fence Agent Components to Configure for clusternode2.example.com"](#) summarizes the components of the HP iLO fence device that this procedure configures for cluster node `clusternode2.example.com`.

Table 5.3. Fence Agent Components to Configure for clusternode2.example.com

Fence Agent Component	Value	Description
Name	hpilofence2	name of the HP iLO fencing device
hostname	hpilohost2	host name of the HP iLO management board to configure as a fence device for clusternode2.example.com
HP iLO login	hpilologin	login identity for the HP iLO management board for clusternode2.example.com
password	hpilopword	password for the HP iLO management board for clusternode2.example.com

[Table 5.4, "Fence Agent Components to Configure for clusternode3.example.com"](#) summarizes the components of the HP iLO fence device that this procedure configures for cluster node `clusternode3.example.com`.

Table 5.4. Fence Agent Components to Configure for clusternode3.example.com

Fence Agent Component	Value	Description
Name	hpilofence3	name of the HP iLO fencing device
hostname	hpilohost3	IP address of the HP iLO management board to configure as a fence device for clusternode3.example.com
HP iLO login	hpilologin	login identity for the HP iLO management board for clusternode3.example.com
password	hpilopword	password for the HP iLO management board for clusternode3.example.com

5.3. HP iLO Fence Device Configuration Procedure

This section provides the procedure for adding an HP iLO fence device to each node of cluster `hpiloclust`. Each node of `hpiloclust` is managed by its own HP iLO management board.

Use the following procedure to configure the HP iLO management board as the fence device for node `clusternode1.example.com` using Conga:

1. As an administrator of `luci` Select the `cluster` tab. This displays the **Choose a cluster to administer** screen.

2. From the **Choose a cluster to administer** screen, you should see the previously configured cluster **hpiclust** displayed, along with the nodes that make up the cluster. Click on **clusternode1.example.com**. This displays the configuration screen for node **clusternode1.example.com**.
3. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
4. From the dropdown menu, under **Create a new Fence Device**, select **HP iLO**. This displays a fence device configuration menu, as shown in [Figure 5.2, "Creating an HP iLO Fence Device"](#).

Main Fencing Method		Backup Fencing Method
Fence Type	HP iLO	Add a fence device to this level
Name	<input type="text"/>	
Hostname	<input type="text"/>	
Login	<input type="text"/>	
Password	<input type="password"/>	
Password Script (optional)	<input type="text"/>	
Use SSH	<input type="checkbox"/>	
<input type="button" value="Remove this device"/>		
Add a fence device to this level		
<input type="button" value="Update main fence properties"/>		<input type="button" value="Update backup fence properties"/>

Figure 5.2. Creating an HP iLO Fence Device

5. For **Name**, enter **hpilofence1**.
6. For **Hostname**, enter **hpilohost1**.
7. For **Login**, enter **hpilologin**.
8. For **Password**, enter **hpilopword**.
9. For **Password Script**, leave the field blank.
10. For **Use SSH**, leave the field blank. You would check this box if your system uses SSH to access the HP iLO management board.
11. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
12. On the confirmation screen, click **OK**. After the fence device has been added, a progress page is displayed after which the display returns to the configuration page for **clusternode1.example.com** in cluster **hpiclust**.

After configuring an HP iLO fence device for **clusternode1.example.com**, use the following procedure to configure an HP iLO fence device for **clusternode2.example.com**.

1. From the configuration page for **clusternode1.example.com**, a menu appears on the left of the screen for cluster **hpiloclust**. Select the node **clusternode2.example.com**. The configuration page for **clusternode2.example.com** appears, with no fence device configured.
2. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
3. From the dropdown menu, under **Create a new Fence Device**, select **HP iLO**. This displays a fence device configuration menu.
4. For **Name**, enter **hpilofence2**.
5. For **Hostname**, enter **hpilohost2**.
6. For **Login**, enter **hpilologin**.
7. For **Password**, enter **hpilopword**.
8. For **Password Script**, leave the field blank.
9. For **Use SSH**, leave the field blank.
10. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
11. On the confirmation screen, click **OK**. After the fence device has been added, a progress page is displayed after which the display returns to the configuration page for **clusternode1.example.com** in cluster **hpiloclust**.

After configuring **hpilofence2** as the fencing device for **clusternode2.example.com**, select node **clusternode3.example.com** from the menu on the left side of the page and configure an HP iLO fence device for that node using the same procedure as you did to configure the fence devices for **clusternode2.example.com** and **clusternode3.example.com**. For **clusternode3.example.com**, use **hpilofence3** as the name of the fencing method and **hpilohost3** as the host name. Otherwise, use the same values for the fence device parameters.

5.4. Cluster Configuration File with HP iLO Fence Device

Configuring a cluster with Conga modifies the cluster configuration file. This section shows the cluster configuration file before and after the procedures documented in [Section 5.3, "HP iLO Fence Device Configuration Procedure"](#) and were performed.

Before the cluster resources and service were configured, the **cluster.conf** file appeared as follows.

```
<?xml version="1.0"?>
<cluster alias="hpiloclust" config_version="12" name="hpiloclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
      <fence/>
    </clusternode>
```

```

        <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
                <fence>
                        <method name="1"/>
                </fence>
        </clusternode>
</clusternodes>
<cman/>
<fencedevices/>
<rm>
        <failoverdomains/>
        <resources/>
</rm>
</cluster>

```

After the cluster resources and service were configured, the `cluster.conf` file appears as follows.

```

<?xml version="1.0"?>
<cluster alias="backupclust" config_version="26" name="backupclust">
  <fence_daemon clean_start="0" post_fail_delay="0" post_join_delay="3"/>
  <clusternodes>
    <clusternode name="doc-10.lab.msp.redhat.com" nodeid="1" votes="1">
      <fence>
        <method name="1">
          <device name="hpilofence1"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="doc-11.lab.msp.redhat.com" nodeid="2" votes="1">
      <fence>
        <method name="1">
          <device name="hpilofence2"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="doc-12.lab.msp.redhat.com" nodeid="3" votes="1">
      <fence>
        <method name="1">
          <device name="hpilofence3"/>
        </method>
      </fence>
    </clusternode>
  </clusternodes>
  <cman/>
  <fencedevices>
    <fencedevice agent="fence_ilo" hostname="hpilohost1"
login="hpilologin" name="hpilofence1" passwd="hpilopword"/>
    <fencedevice agent="fence_ilo" hostname="hpilohost2"
login="hpilologin" name="hpilofence2" passwd="hpilologin"/>
    <fencedevice agent="fence_ilo" hostname="hpilohost3"
login="hpilologin" name="hpilofence3" passwd="hpilopword"/>
  </fencedevices>
  <rm>
    <failoverdomains/>
  </rm>
</cluster>

```

```
        <resources/>
    </rm>
</cluster>
```

5.5. Testing the HP iLO Fence Device Configuration

To check whether the configuration you have defined works as expected, you can use the **fence_node** to fence a node manually. The **fence_node** program reads the fencing settings from the **cluster.conf** file for the given node and then runs the configured fencing agent against the node.

To test whether the HP iLO management boards have been successfully configured as fence devices for the three nodes in cluster **hpioloclust**, execute the following commands and check whether the nodes have been fenced.

```
# /sbin/fence_node clusternode1.example.com
# /sbin/fence_node clusternode2.example.com
# /sbin/fence_node clusternode3.example.com
```

Chapter 6. Configuring Fencing with Dual Power Supplies

If your system is configured with redundant power supplies for your system, you must be sure to configure fencing so that your nodes fully shut down when they need to be fenced. If you configure each power supply as a separate fence method, each power supply will be fenced separately; the second power supply will allow the system to continue running when the first power supply is fenced and the system will not be fenced at all. To configure a system with dual power supplies, you must configure your fence devices so that both power supplies are shut off and the system is taken completely down. This requires that you configure two fencing devices inside of a single fencing method.

This chapter provides the procedures for using the Conga configuration tool in a Red Hat cluster to configure fencing with dual power supplies.

[Figure 6.1, “Fence Devices with Dual Power Supplies”](#) shows the configuration this procedure yields. In this configuration, there are two APC network power switches, each of which runs on its own separate UPS and has its own unique IP address. Each node in the cluster is connected to a port on each APC switch.

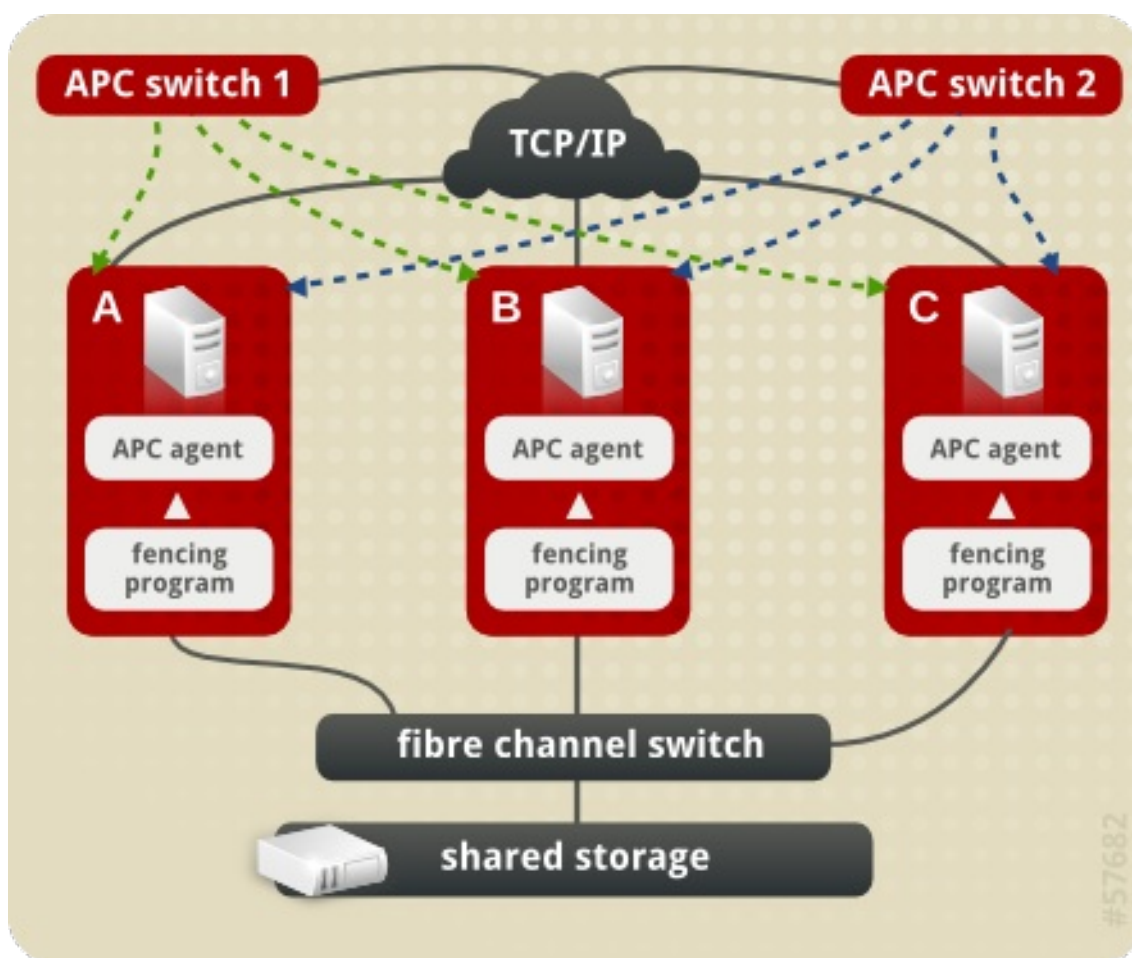


Figure 6.1. Fence Devices with Dual Power Supplies

6.1. Dual Power Fencing Prerequisite Configuration

[Table 6.1, “Configuration Prerequisites”](#) summarizes the prerequisite components that have been set up before this procedure begins.

Table 6.1. Configuration Prerequisites

Component	Name	Comment
cluster	apcclust	three-node cluster
cluster node	clusternode1.example.com	node in cluster apcclust configured with 2 APC switches to administer power supply
cluster node	clusternode2.example.com	node in cluster apcclust configured with 2 APC switches to administer power supply
cluster node	clusternode3.example.com	node in cluster apcclust configured with 2 APC switches to administer power supply
IP address	10.15.86.96	IP address for the first APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com. This switch runs on its own UPS.
IP address	10.15.86.97	IP address for the second APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com. This switch runs on its own UPS.

[Table 6.2, "Configuration Prerequisites"](#) summarizes the prerequisite components that have been set up for each of the APC switches before this procedure begins.

Table 6.2. Configuration Prerequisites

Component	Name	Comment
login	apclogin	login value for both of the the APC switches that control the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
password	apcpword	password for both the APC switches that control the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
port	1	port number on both of the APC switches that clusternode1.example.com connects to
port	2	port number on both of the APC switches that clusternode2.example.com connects to
port	3	port number on both of the APC switches that clusternode3.example.com connects to

6.2. Fence Device Components to Configure

This procedure configures two APC switches as fence devices that will be used for each node in cluster **apcclust**. Then the procedure configures both of those switches as part of one fencing method for **clusternode1.example.com**, **clusternode2.example.com**, and **clusternode1.example.com**.

[Table 6.3, "Fence Device Components to Configure for APC Fence Device"](#) summarizes the components of the APC fence devices that this procedure configures for cluster node **clusternode1.example.com**.

Table 6.3. Fence Device Components to Configure for APC Fence Device

Fence Device Component	Value	Description
Fencing Type	APC Power Switch	type of fencing device to configure for each APC switch

Fence Device Component	Value	Description
Name	pwr01	name of the first APC fencing device for node1.example.com, node2.example.com, and node3.example.com
IP address	10.15.86.96	IP address of the first APC switch to configure as a fence device for node1.example.com, node2.example.com, and node3.example.com
Name	pwr02	name of the second APC fencing device for node1.example.com, node2.example.com, and node3.example.com
IP address	10.15.86.97	IP address of the second APC switch to configure as a fence device for node1.example.com, node2.example.com, and node3.example.com
login	apclogin	login value for the each of the APC switches that control the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
password	apcpassword	password for each of the APC switches that control the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com

[Table 6.4, “Fence Agent Components to Specify for Each Node in apcclust”](#) summarizes the components of each of the APC fence devices that you must specify for the cluster nodes in **apcclust**.

Table 6.4. Fence Agent Components to Specify for Each Node in apcclust

Fence Agent Component	Value	Description
fence device	pwr01	name of the first APC fence device you defined as a shared device
fence device	pwr02	name of the second APC fence device you defined as a shared device
port	1	port number on each of the APC switches for node1.example.com
port	2	port number on each of the APC switches for node2.example.com
port	3	port number on each of the APC switches for node3.example.com

The remainder of the fence device components that you configure for each fence device for each node appear automatically when you specify that you will be configuring the **pwr01** or **pwr02** fence device that you previously defined as a shared fence device.

6.3. Dual Power Fencing Configuration Procedure

This section provides the procedure for adding two APC fence devices to each node of cluster **apcclust**, configured as a single fence method to ensure that the fencing is successful. This example uses the same APC switches for each cluster node. The APC switches will first be configured as shared fence devices. After configuring the APC switches as shared fence devices, the devices will be added as fence device for each node in the cluster.

To configure the first APC switch as a shared fence device named **pwr01** using **Conga**, perform the following procedure:

1. As an administrator of **luci** Select the **cluster** tab. This displays the **Choose a cluster to administer** screen.

- From the **Choose a cluster to administer** screen, you should see the previously configured cluster **apcclust** displayed, along with the nodes that make up the cluster. Click on **apcclust** to select the cluster.
- At the detailed menu for the cluster **apcclust** (below the **clusters** menu on the left side of the screen), click **Shared Fence Devices**. Clicking **Shared Fence Devices** causes the display of any shared fence devices previously configured for a cluster and causes the display of menu items for fence device configuration: **Add a Fence Device** and **Configure a Fence Device**.
- Click **Add a Fence Device**. Clicking **Add a Fence Device** causes the **Add a Sharable Fence Device** page to be displayed.
- At the **Add a Sharable Fence Device** page, click the drop-down box under **Fencing Type** and select **APC Power Switch**. This causes Conga to display the components of an APC Power Switch fencing type, as shown in [Figure 6.2, "Adding a Sharable Fence Device"](#).

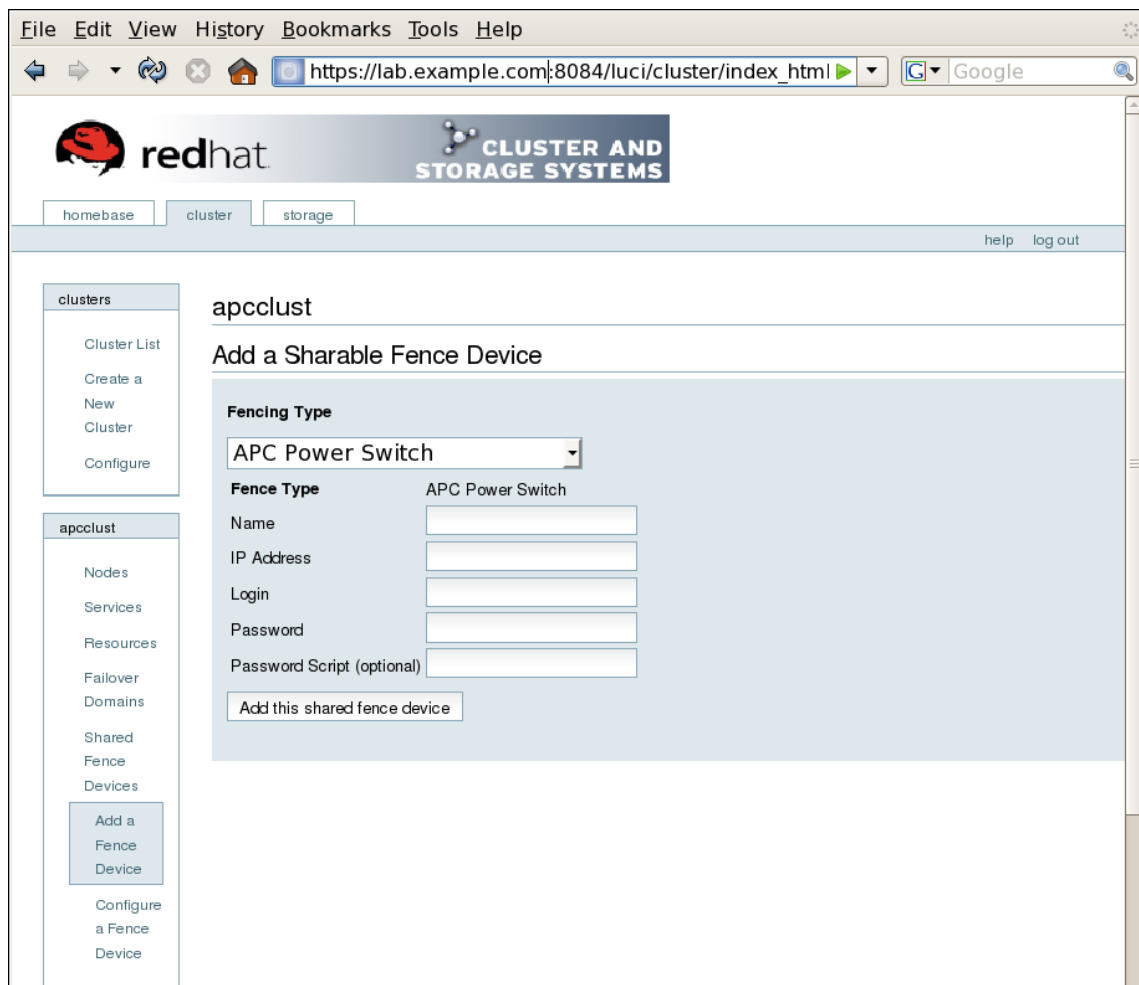


Figure 6.2. Adding a Sharable Fence Device

- For **Name**, enter **pwr01**.
- For **IP Address**, enter **10.15.86.96**.
- For **Login**, enter **apclogin**.
- For **Password**, enter **apcpassword**.
- For **Password Script**, leave blank.

11. Click **Add this shared fence device**.

Clicking **Add this shared fence device** causes a progress page to be displayed temporarily. After the fence device has been added, the detailed cluster properties menu is updated with the fence device under **Configure a Fence Device**.

To configure the second APC switch as a shared fence device named **pwr02**, perform the following procedure:

1. After configuring the first APC switch as shared fence device **pwr01**, click **Add a Fence Device** from the detailed menu for the cluster **apcclust** (below the **clusters** menu on the left side of the screen). This displays the **Add a Sharable Fence Device** page.
2. At the **Add a Sharable Fence Device** page, click the drop-down box under **Fencing Type** and select **APC Power Switch**. This causes Conga to display the components of an APC Power Switch fencing type.
3. For **Name**, enter **pwr02**.
4. For **IP Address**, enter **10.15.86.97**.
5. For **Login**, enter **apclogin**.
6. For **Password**, enter **apcpword**.
7. For **Password Script**, leave blank.
8. Click **Add this shared fence device**.

Clicking **Add this shared fence device** causes a progress page to be displayed temporarily. After the fence device has been added, the detailed cluster properties menu is updated with the fence device under **Configure a Fence Device**.

After configuring the APC switches as shared fence devices, use the following procedure to configure the first APC switch, **pwr01**, as the first fence device for node **clusternode1.example.com**.

1. At the detailed menu for the cluster **apcclust** (below the **clusters** menu), click **Nodes**. Clicking **Nodes** causes the display of the status of each node in **apcclust**.
2. At the bottom of the display for node **clusternode1.example.com**, click **Manage Fencing for this Node**. This displays the configuration screen for node **clusternode1.example.com**.
3. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
4. From the dropdown menu, the **pwr01** and **pwr02** fence devices you have already created should display as one of the menu options under **Use an Existing Fence Device**. Select **pwr01 (APC Power Device)**. This causes a fence device configuration menu to display with the **Name**, **IP Address**, **Login**, **Password**, and **Password Script** values already configured, as defined when you configured **pwr01** as a shared fence device. (The **Password** value does not display, but you may not alter it.) This is shown in [Figure 6.3, "Adding Fence Device pwr01 to a Node"](#).

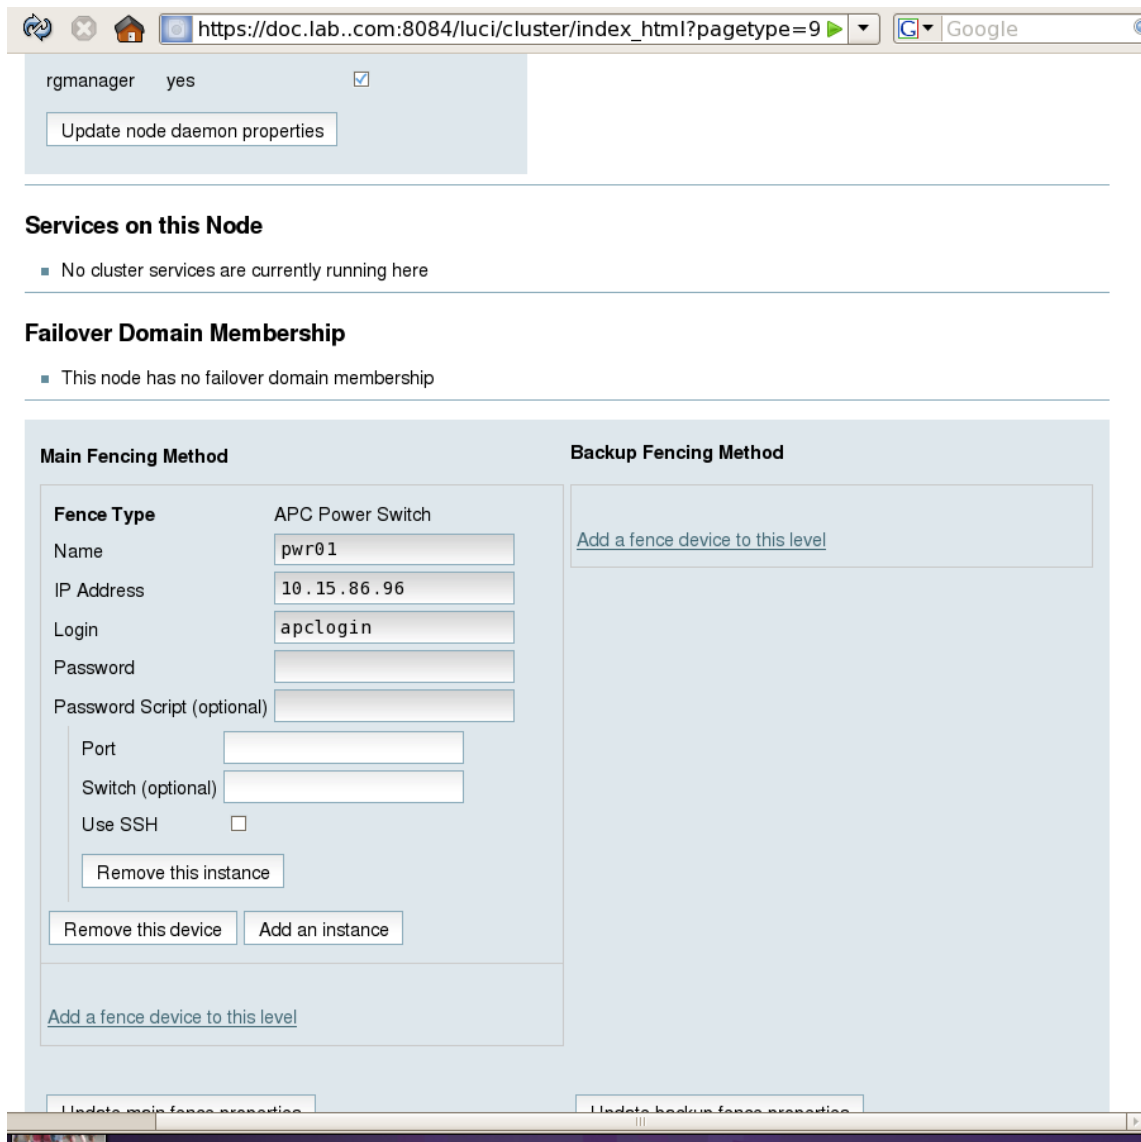


Figure 6.3. Adding Fence Device pwr01 to a Node

5. For **Port**, enter **1**. Do not enter any value for **Switch**.

Before updating the main fence properties for this node, use the following procedure to add **pwr02** as the second fence device of the main fencing method for node **clusternode1.example.com**.

1. Beneath the configuration information for **pwr01** that you have entered, click **Add a fence device to this level**. This displays the dropdown menu again.
2. From the dropdown menu, select **pwr02 (APC Power Device)**. This causes a fence device configuration menu to display with the **Name**, **IP Address**, **Login**, **Password**, and **Password Script** values already configured, as defined when you configured **pwr02** as a shared fence device. This is shown in [Figure 6.4, “Adding Fence Device pwr02 to a Node”](#).

Main Fencing Method	Backup Fencing Method
<p>Fence Type APC Power Switch</p> <p>Name <input type="text" value="pwr01"/></p> <p>IP Address <input type="text" value="10.15.86.96"/></p> <p>Login <input type="text" value="apclogin"/></p> <p>Password <input type="password"/></p> <p>Password Script (optional) <input type="text"/></p> <p>Port <input type="text" value="1"/></p> <p>Switch (optional) <input type="text"/></p> <p>Use SSH <input type="checkbox"/></p> <p><input type="button" value="Remove this instance"/></p> <p><input type="button" value="Remove this device"/> <input type="button" value="Add an instance"/></p>	<p>Add a fence device to this level</p>
<p>Fence Type APC Power Switch</p> <p>Name <input type="text" value="pwr02"/></p> <p>IP Address <input type="text" value="10.15.86.97"/></p> <p>Login <input type="text" value="apclogin"/></p> <p>Password <input type="password"/></p> <p>Password Script (optional) <input type="text"/></p> <p>Port <input type="text"/></p> <p>Switch (optional) <input type="text"/></p> <p>Use SSH <input type="checkbox"/></p> <p><input type="button" value="Remove this instance"/></p> <p><input type="button" value="Remove this device"/> <input type="button" value="Add an instance"/></p>	

Figure 6.4. Adding Fence Device pwr02 to a Node

3. For **Port**, enter **1**. Do not enter any value for **Switch**.

After entering the configuration information for both power sources to use as fence devices, you can update the main fence properties using the following procedure.

1. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
2. On the confirmation screen, Click **OK**. A progress page is displayed after which the display returns to the status page for **clusternode1.example.com** in cluster **apcclust**.

After configuring **pwr01** and **pwr02** as the fencing devices for **clusternode1.example.com**, use the same procedure to configure these same devices as the fencing devices for **clusternode2.example.com**, specifying Port 2 on each switch for **clusternode2.example.com**:

1. On the status page for **clusternode1.example.com** in cluster **apcclust**, the other nodes in **apcclust** are displayed below the **Configure** menu item below the **Nodes** menu item on the left side of the screen. Click **clusternode2.example.com** to display the status screen for **clusternode2.example.com**.
2. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
3. As for **clusternode1.example.com**, the fence device **pwr01** should display as one of the menu options on the dropdown menu, under **Use an Existing Fence Device**. Select **pwr01 (APC Power Device)**. This causes a fence device configuration menu to display with the **Name, IP Address, Login, Password, Password Script** values already configured, as defined when you configured **pwr01** as a shared fence device.
4. For **Port**, enter **2**. Do not enter any value for **Switch**.
5. Before clicking on **Update main fence properties**, click on **Add a fence device to this level** to add the fence device **pwr02**.
6. Select **pwr02 (APC Power Device)** from the **Use an Existing Fence Device** display of the dropdown menu. This causes a fence device configuration menu to display with the **Name, IP Address, Login, Password, Password Script** values already configured, as defined when you configured **pwr01** as a shared fence device.
7. For **Port**, enter **2**. Do not enter any value for **Switch**.
8. To configure both of the fence devices, Click **Update main fence properties**.

Similarly, configure **pwr01** and **pwr02** as the main fencing method for **clusternode3.example.com**, this time specifying **3** as the Port number for both devices.

6.4. Cluster Configuration File with Dual Power Supply Fencing

Configuring a cluster with Conga modifies the cluster configuration file. This section shows the cluster configuration file before and after the procedures documented in [Section 6.3, "Dual Power Fencing Configuration Procedure"](#) were performed.

Before the cluster resources and service were configured, the **cluster.conf** file appeared as follows.

```
<?xml version="1.0"?>
<cluster alias="apcclust" config_version="34" name="apcclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
      <fence/>
```

```

        </clusternode>
        <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
                </fence>
        </clusternode>
</clusternodes>
<cman/>
<fencedevices/>
<rm>
        <failoverdomains/>
        <resources/>
</rm>
</cluster>

```

After the cluster resources and service were configured, the `cluster.conf` file appears as follows.

```

<?xml version="1.0"?>
<cluster alias="apcclust" config_version="40" name="apcclust">
  <fence_daemon clean_start="0" post_fail_delay="0" post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1" votes="1">
      <fence>
        <method name="1">
          <device name="pwr01" option="off"
port="1"/>
          <device name="pwr02" option="off"
port="1"/>
          <device name="pwr01" option="on"
port="1"/>
          <device name="pwr02" option="on"
port="1"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2" votes="1">
      <fence>
        <method name="1">
          <device name="pwr01" option="off"
port="2"/>
          <device name="pwr02" option="off"
port="2"/>
          <device name="pwr01" option="on"
port="2"/>
          <device name="pwr02" option="on"
port="2"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="clusternode3.example.com" nodeid="3" votes="1">
      <fence>
        <method name="1">
          <device name="pwr01" option="off"
port="3"/>
          <device name="pwr02" option="off"

```

```

port="3"/>
                                <device name="pwr01" option="on"
port="3"/>
                                <device name="pwr02" option="on"
port="3"/>
                                </method>
                                </fence>
                                </clusternode>
                                </clusternodes>
                                <cman/>
                                <fencedevices>
                                <fencedevice agent="fence_apc" ipaddr="10.15.86.96"
login="apclogin" name="pwr01" passwd="apcpword"/>
                                <fencedevice agent="fence_apc" ipaddr="10.15.86.97"
login="apclogin" name="pwr02" passwd="apcpword"/>
                                </fencedevices>
                                <rm>
                                <failoverdomains/>
                                <resources/>
                                </rm>
                                </cluster>

```

6.5. Testing the Dual Power Fence Device Configuration

To check whether the configuration you have defined works as expected, you can use the **fence_node** to fence a node manually. The **fence_node** program reads the fencing settings from the **cluster.conf** file for the given node and then runs the configured fencing agent against the node.

To test whether the dual power fencing configuration been successfully configured for the three nodes in cluster **apcclust**, execute the following commands and check whether the nodes have been fenced.

```

# /sbin/fence_node clusternode1.example.com
# /sbin/fence_node clusternode2.example.com
# /sbin/fence_node clusternode3.example.com

```


Chapter 7. Configuring a Backup Fencing Method

You can define multiple fencing methods for a node. If fencing fails using the first method, the system will attempt to fence the node using the second method. This chapter provides the procedures for using the Conga configuration tool in a Red Hat cluster to configure a main fencing method and a backup fencing method.

[Figure 7.1, “Cluster Configured with Backup Fencing Method”](#) shows the configuration this procedure yields. In this configuration, the main fencing method consists of two APC network power switches, each of which runs on its own separate UPS and has its own unique IP address. Each node in the cluster is connected to a port on each APC switch. As a backup fencing method, each node on this cluster is configured with an IPMI management board as a fencing device.



Note

Note that in this configuration each system has redundant power and is hooked into two independent power sources. This ensures that the IPMI management board in the node would still function as needed in a cluster even if you lose power from one of the sources.

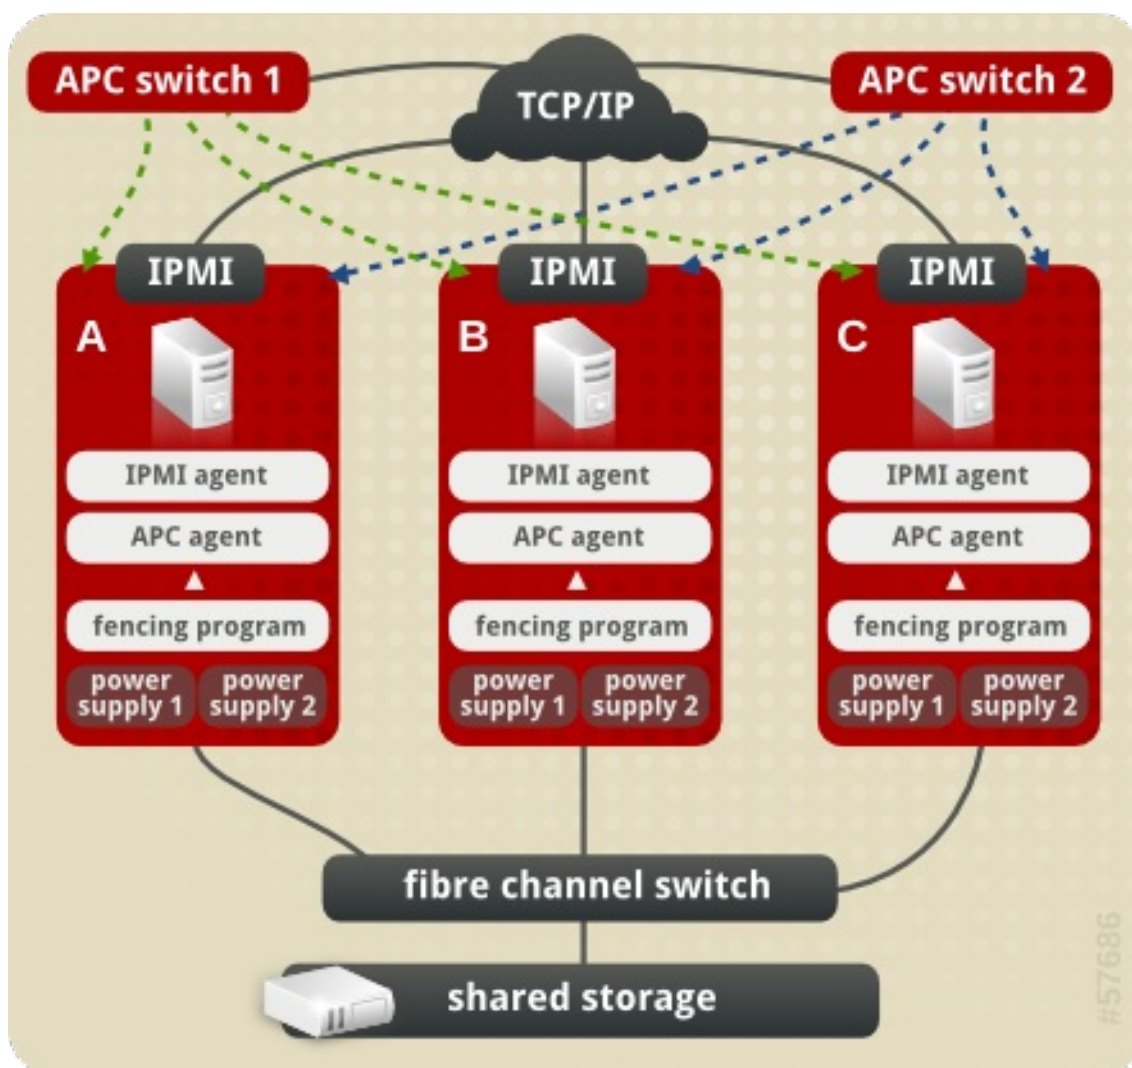


Figure 7.1. Cluster Configured with Backup Fencing Method

7.1. Backup Fencing Prerequisite Configuration

[Table 7.1, "Configuration Prerequisites"](#) summarizes the prerequisite components that have been set up before this procedure begins.

Table 7.1. Configuration Prerequisites

Component	Name	Comment
cluster	backupclust	three-node cluster
cluster node	clusternode1.example.com	node in cluster backupclust configured with 2 APC switches, an IPMI management board, and 2 power supplies
cluster node	clusternode2.example.com	node in cluster backupclust configured with 2 APC switches, an IPMI management board, and 2 power supplies
cluster node	clusternode3.example.com	node in cluster backupclust configured with 2 APC switches, an IPMI management board, and 2 power supplies
IP address	10.15.86.96	IP address for the first APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com. This switch runs on its own UPS.
IP address	10.15.86.97	IP address for the second APC switch that controls the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com. This switch runs on its own UPS.
IP address	10.15.86.50	IP address for IPMI management board for clusternode1.example.com
IP address	10.15.86.51	IP address for IPMI management board for clusternode2.example.com
IP address	10.15.86.52	IP address for IPMI management board for clusternode3.example.com

[Table 7.2, "Configuration Prerequisites"](#) summarizes the prerequisite components that have been set up for each of the APC switches before this procedure begins.

Table 7.2. Configuration Prerequisites

Component	Name	Comment
login	apclogin	login value for both of the the APC switches that control the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
password	apcpword	password for both the APC switches that control the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
port	1	port number on both of the APC switches that clusternode1.example.com connects to
port	2	port number on both of the APC switches that clusternode2.example.com connects to
port	3	port number on both of the APC switches that clusternode3.example.com connects to

[Table 7.3, "Configuration Prerequisites"](#) summarizes the prerequisite components that have been set up for each of the IPMI management boards before this procedure begins.

Table 7.3. Configuration Prerequisites

Component	Name	Comment
login	ipmilogin	login name for IPMI management board for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com
password	ipmipword	password IPMI management board for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com

7.2. Fence Device Components to Configure

This procedure consists of the following steps:

1. The procedure configures two APC switches as fence devices that will be used as the main fencing method for each node in cluster **backupclust**.
2. The procedure configures the main and backup fencing methods for **clusternode1.example.com**, using the two APC switches for the main fencing method for the node and using its IPMI management board as the backup fencing method for the node.
3. The procedure configures the main and backup fencing methods for **clusternode2.example.com**, using the two APC switches for the main fencing method for the node and using its IPMI management board as the backup fencing method for the node.
4. The procedure configures the main and backup fencing methods for **clusternode3.example.com**, using the two APC switches for the main fencing method for the node and using its IPMI management board as the backup fencing method for the node.

[Table 7.4, “Fence Device Components to Configure for APC Fence Device”](#) summarizes the components of the APC fence devices that this procedure configures for the nodes in the cluster **backupclust**.

Table 7.4. Fence Device Components to Configure for APC Fence Device

Fence Device Component	Value	Description
Fencing Type	APC Power Switch	type of fencing device to configure for each APC switch
Name	pwr01	name of the first APC fencing device for node1.example.com, node2.example.com, and node3.example.com
IP address	10.15.86.96	IP address of the first APC switch to configure as a fence device for node1.example.com, node2.example.com, and node3.example.com
Name	pwr02	name of the second APC fencing device for node1.example.com, node2.example.com, and node3.example.com
IP address	10.15.86.97	IP address of the second APC switch to configure as a fence device for node1.example.com, node2.example.com, and node3.example.com
login	apclogin	login value for the each of the APC switches that control the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com

Fence Device Component	Value	Description
password	apcpword	password for each of the APC switches that control the power for for clusternode1.example.com, clusternode2.example.com, and clusternode3.example.com

[Table 7.5, “Fence Agent Components to Specify for clusternode1.example.com”](#) summarizes the components of the main and backup fence devices that you specify for the node **clusternode1.example.com**.

Table 7.5. Fence Agent Components to Specify for clusternode1.example.com

Fence Agent Component	Value	Description
fence device	pwr01	name of the first APC fence device you defined as a shared device
port	1	port number on the first APC switch for node1.example.com
fence device	pwr02	name of the second APC fence device you defined as a shared device
port	1	port number on the second APC switch for clusternode1.example.com
Name	ipmifence1	name of the IPMI fencing device for clusternode1.example.com
IP address	10.15.86.50	IP address of the IPMI management board for clusternode1.example.com
IPMI login	ipmilogin	login identity for the IPMI management board for clusternode1.example.com
password	ipmipword	password for the IPMI management board for clusternode1.example.com
authentication type	password	authentication type for the IPMI management board for clusternode1.example.com

[Table 7.6, “Fence Agent Components to Specify for clusternode2.example.com”](#) summarizes the components of the main and backup fence devices that you specify for the node **clusternode2.example.com**.

Table 7.6. Fence Agent Components to Specify for clusternode2.example.com

Fence Agent Component	Value	Description
fence device	pwr01	name of the first APC fence device you defined as a shared device
port	2	port number on the first APC switch for node2.example.com
fence device	pwr02	name of the second APC fence device you defined as a shared device
port	2	port number on the second APC switch for clusternode2.example.com
Name	ipmifence2	name of the IPMI fencing device for clusternode2.example.com
IP address	10.15.86.51	IP address of the IPMI management board for clusternode2.example.com
IPMI login	ipmilogin	login identity for the IPMI management board for clusternode2.example.com
password	ipmipword	password for the IPMI management board for clusternode2.example.com
authentication type	password	authentication type for the IPMI management board for clusternode2.example.com

Table 7.7, “Fence Agent Components to Specify for `clusternode3.example.com`” summarizes the components of the main and backup fence devices that you specify for the node `clusternode3.example.com`.

Table 7.7. Fence Agent Components to Specify for `clusternode3.example.com`

Fence Agent Component	Value	Description
fence device	pwr01	name of the first APC fence device you defined as a shared device
port	3	port number on the first APC switch for <code>node3.example.com</code>
fence device	pwr02	name of the second APC fence device you defined as a shared device
port	3	port number on the second APC switch for <code>clusternode3.example.com</code>
Name	ipmifence3	name of the IPMI fencing device for <code>clusternode3.example.com</code>
IP address	10.15.86.52	IP address of the IPMI management board for <code>clusternode3.example.com</code>
IPMI login	ipmilogin	login identity for the IPMI management board for <code>clusternode3.example.com</code>
password	ipmipword	password for the IPMI management board for <code>clusternode3.example.com</code>
authentication type	password	authentication type for the IPMI management board for <code>clusternode3.example.com</code>

7.3. Backup Fencing Configuration Procedure

This section provides the procedure for adding two APC fence devices to each node of cluster `backupclust`, configured as a single main fence method to ensure that the fencing is successful. This procedure also configures an IPMI management board as a backup fence device for each node of cluster `backupclust`.

This example uses the same APC switches for each cluster node. The APC switches will first be configured as shared fence devices. After configuring the APC switches as shared fence devices, the APC devices and the IPMI devices will be added as fence devices for each node in the cluster.

7.3.1. Configuring the APC switches as shared fence devices

To configure the first APC switch as a shared fence device named `pwr01` using **Conga**, perform the following procedure:

1. As an administrator of **luci** Select the `cluster` tab. This displays the **Choose a cluster to administer** screen.
2. From the **Choose a cluster to administer** screen, you should see the previously configured cluster `backupclust` displayed, along with the nodes that make up the cluster. Click on `backupclust` to select the cluster.
3. At the detailed menu for the cluster `backupclust` (below the **clusters** menu on the left side of the screen), click **Shared Fence Devices**. Clicking **Shared Fence Devices** causes the display of any shared fence devices previously configured for a cluster and causes the display of menu items for fence device configuration: **Add a Fence Device** and **Configure a Fence Device**.
4. Click **Add a Fence Device**. Clicking **Add a Fence Device** causes the **Add a Sharable Fence Device** page to be displayed.

- At the **Add a Sharable Fence Device** page, click the drop-down box under **Fencing Type** and select **APC Power Switch**. This causes Conga to display the components of an APC Power Switch fencing type, as shown in [Figure 3.2, "Adding a Sharable Fence Device"](#).

The screenshot shows a web browser window with the URL `https://lab.example.com:8084/luci/cluster/index_html`. The page title is "Luci — cluster — fence devices - Add a new fence device - Mozilla Firefox". The interface includes a navigation menu with "homebase", "cluster", and "storage" tabs. A sidebar on the left shows a tree view for "clusters" and "backupclust", with "Add a Fence Device" selected. The main content area is titled "backupclust" and "Add a Sharable Fence Device". The "Fencing Type" dropdown is set to "APC Power Switch". Below this, the "Fence Type" is also "APC Power Switch". The form contains five input fields: "Name", "IP Address", "Login", "Password", and "Password Script (optional)". A button labeled "Add this shared fence device" is at the bottom of the form.

Figure 7.2. Adding a Sharable Fence Device

- For **Name**, enter **pwr01**.
- For **IP Address**, enter **10.15.86.96**.
- For **Login**, enter **apclogin**.
- For **Password**, enter **apcpword**.
- For **Password Script**, leave blank.
- Click **Add this shared fence device**.

Clicking **Add this shared fence device** temporarily displays a progress page. After the fence device has been added, the detailed cluster properties menu is updated with the fence device under **Configure a Fence Device**.

To configure the second APC switch as a shared fence device named **pwr02**, perform the following procedure:

1. After configuring the first APC switch as shared fence device **pwr01**, click **Add a Fence Device** from the detailed menu for the cluster **backupclust** (below the **clusters** menu on the left side of the screen). This displays the **Add a Sharable Fence Device** page.
2. At the **Add a Sharable Fence Device** page, click the drop-down box under **Fencing Type** and select **APC Power Switch**. This causes Conga to display the components of an APC Power Switch fencing type.
3. For **Name**, enter **pwr02**.
4. For **IP Address**, enter **10.15.86.97**.
5. For **Login**, enter **apclogin**.
6. For **Password**, enter **apcpword**.
7. For **Password Script**, leave blank.
8. Click **Add this shared fence device**.

Clicking **Add this shared fence device** causes a progress page to be displayed temporarily. After the fence device has been added, the detailed cluster properties menu is updated with the fence device under **Configure a Fence Device**.

7.3.2. Configuring Fencing on the First Cluster Node

After configuring the APC switches as shared fence devices, use the following procedure to configure the first APC switch, **pwr01**, as the first fence device for node **clusternode1.example.com**.

1. At the detailed menu for the cluster **backupclust** (below the **clusters** menu), click **Nodes**. Clicking **Nodes** causes the display of the status of each node in **backupclust**.
2. At the bottom of the display for node **clusternode1.example.com**, click **Manage Fencing for this Node**. This displays the configuration screen for node **clusternode1.example.com**.
3. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
4. From the dropdown menu, the **pwr01** and **pwr02** fence devices you have already created should display as one of the menu options under **Use an Existing Fence Device**. Select **pwr01 (APC Power Device)**. This causes a fence device configuration menu to display with the **Name**, **IP Address**, **Login**, **Password**, and **Password Script** values already configured, as defined when you configured **pwr01** as a shared fence device. (The **Password** value does not display, but you may not alter it.) This is shown in [Figure 7.3, "Adding Fence Device pwr01 to a Node"](#).

Main Fencing Method	Backup Fencing Method
<p>Fence Type APC Power Switch</p> <p>Name <input type="text" value="pwr01"/></p> <p>IP Address <input type="text" value="10.15.86.96"/></p> <p>Login <input type="text" value="apclogin"/></p> <p>Password <input type="password"/></p> <p>Password Script (optional) <input type="text"/></p> <p>Port <input type="text" value="1"/></p> <p>Switch (optional) <input type="text"/></p> <p>Use SSH <input type="checkbox"/></p> <p><input type="button" value="Remove this instance"/></p> <p><input type="button" value="Remove this device"/> <input type="button" value="Add an instance"/></p> <p>Add a fence device to this level</p> <p><input type="button" value="Update main fence properties"/></p>	<p>Add a fence device to this level</p> <p><input type="button" value="Update backup fence properties"/></p>

Figure 7.3. Adding Fence Device pwr01 to a Node

5. For **Port**, enter **1**. Do not enter any value for **Switch**.

Before updating the main fence properties for this node, use the following procedure to add **pwr02** as the second fence device of the main fencing method for node **clusternode1.example.com**.

1. Beneath the configuration information for **pwr01** that you have entered, click **Add a fence device to this level**. This displays the dropdown menu again.
2. From the dropdown menu, select **pwr02 (APC Power Device)**. This causes a fence device configuration menu to display with the **Name**, **IP Address**, **Login**, **Password**, and **Password Script** values already configured, as defined when you configured **pwr02** as a shared fence device. This is shown in [Figure 7.4, "Adding Fence Device pwr02 to a Node"](#).

The screenshot displays two identical configuration panels for fence devices. The top panel is for device 'pwr01' with IP '10.15.86.96'. The bottom panel is for device 'pwr02' with IP '10.15.86.97'. Both are configured as 'APC Power Switch' devices with login 'apclogin'. The 'Port' field in the top panel is set to '1', while the bottom panel's 'Port' field is empty. A 'Switch (optional)' field is present in both but empty. A 'Use SSH' checkbox is unchecked in both. Each panel includes a 'Remove this instance' button and a shared 'Remove this device' and 'Add an instance' button at the bottom. A link 'Add a fence device to this level' is located in the top right corner of the interface.

Figure 7.4. Adding Fence Device pwr02 to a Node

3. For **Port**, enter **1**. Do not enter any value for **Switch**.

After entering the configuration information for both power sources to use as fence devices, you can update the main fence properties using the following procedure.

1. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
2. On the confirmation screen, Click **OK**. A progress page is displayed after which the display returns to the status page for **clusternode2.example.com** in cluster **backupclust**.

After configuring the main fence method for **clusternode1.example.com** and updating the main fence properties, use the following procedure to configure the IPMI management board for node **clusternode1.example.com** as the backup fencing method for that node:

1. At the **Backup Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
2. From the dropdown menu, under **Create a new Fence Device**, select **IPMI Lan**. This displays a fence device configuration menu, as shown in [Figure 7.5, "Configuring a Backup Fencing Method"](#).

Main Fencing Method	Backup Fencing Method
<p>Fence Type APC Power Switch</p> <p>Name <input type="text" value="pwr01"/></p> <p>IP Address <input type="text" value="10.15.86.96"/></p> <p>Login <input type="text" value="apclogin"/></p> <p>Password <input type="password"/></p> <p>Password Script (optional) <input type="text"/></p> <p>Port <input type="text" value="1"/></p> <p>Switch (optional) <input type="text"/></p> <p>Use SSH <input type="checkbox"/></p> <p><input type="button" value="Remove this instance"/></p> <p><input type="button" value="Remove this device"/> <input type="button" value="Add an instance"/></p>	<p>Fence Type IPMI Lan</p> <p>Name <input type="text"/></p> <p>IP Address <input type="text"/></p> <p>Login <input type="text"/></p> <p>Password <input type="password"/></p> <p>Password Script (optional) <input type="text"/></p> <p>Authentication Type <input type="text"/></p> <p>Use Lanplus <input type="checkbox"/></p> <p><input type="button" value="Remove this device"/></p> <p>Add a fence device to this level</p>
<p>Fence Type APC Power Switch</p> <p>Name <input type="text" value="pwr02"/></p> <p>IP Address <input type="text" value="10.15.86.97"/></p> <p>Login <input type="text" value="apclogin"/></p> <p>Password <input type="password"/></p> <p>Password Script (optional) <input type="text"/></p> <p>Port <input type="text" value="1"/></p> <p>Switch (optional) <input type="text"/></p>	

Figure 7.5. Configuring a Backup Fencing Method

3. For **Name**, enter **ipmifence1**.
4. For **IP Address**, enter **10.15.86.50**.
5. For **Login**, enter **ipmilogin**.
6. For **Password**, enter **ipmipword**.
7. For **Password Script**, leave the field blank.
8. For **Authentication type**, enter **password**. This field specifies the IPMI authentication type. Possible values for this field are none, **password**, **md2**, or **md5**.
9. Leave the **Use Lanplus** field blank. You would check this field if your fence device is a Lanplus-capable interface such as iLO2.

After entering the configuration information for the backup fencing method for **clusternode1.example.com**, you can update the backup fence properties using the following procedure.

1. Click **Update backup fence properties** at the bottom of the right side of the screen. This causes a confirmation screen to be displayed.
2. On the confirmation screen, click **OK**. After the fence device has been added, a progress page is displayed after which the display returns to the configuration page for **clusternode1.example.com** in cluster **backupclust**.

7.3.3. Configuring Fencing on the Remaining Cluster Nodes

After configuring the main fencing method and the backup fencing method for **clusternode1.example.com**, use the same procedure to configure the fencing methods for **clusternode2.example.com** and **clusternode3.example.com**.

1. At the detailed menu for the cluster **backupclust** (below the **clusters** menu on the left side of the screen) click on **clusternode2.example.com**, which should be displayed below **Nodes** -> **Configure**. This displays the configuration screen for node **clusternode2.example.com**.
2. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
3. From the dropdown menu, the **pwr01** and **pwr02** fence devices you have already created should display as one of the menu options under **Use an Existing Fence Device**. Select **pwr01 (APC Power Device)**. This causes a fence device configuration menu to display with the **Name**, **IP Address**, **Login**, **Password**, and **Password Script** values already configured, as defined when you configured **pwr01** as a shared fence device. (The **Password** value does not display, but you may not alter it.)
4. For **Port**, enter **2**. Do not enter any value for **Switch**.

Before updating the main fence properties for this node, use the following procedure to add **pwr02** as the second fence device of the main fencing method for node **clusternode1.example.com**.

1. Beneath the configuration information for **pwr01** that you have entered, click **Add a fence device to this level**. This displays the dropdown menu again.
2. From the dropdown menu, select **pwr02 (APC Power Device)**. This causes a fence device configuration menu to display with the **Name**, **IP Address**, **Login**, **Password**, and **Password Script** values already configured, as defined when you configured **pwr02** as a shared fence device.
3. For **Port**, enter **2**. Do not enter any value for **Switch**.

After entering the configuration information for both power sources to use as fence devices, you can update the main fence properties using the following procedure.

1. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
2. On the confirmation screen, Click **OK**. A progress page is displayed after which the display returns to the status page for **clusternode1.example.com** in cluster **backupclust**.

After configuring the main fence method for **clusternode2.example.com** and updating the main fence properties, use the following procedure to configure the IPMI management board for node **clusternode2.example.com** as the backup fencing method for that node:

1. At the **Backup Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
2. From the dropdown menu, under **Create a new Fence Device**, select **IPMI Lan**. This displays a fence device configuration menu.
3. For **Name**, enter **ipmifence1**.
4. For **IP Address**, enter **10.15.86.51**.
5. For **Login**, enter **ipmilogin**.
6. For **Password**, enter **ipmipword**.
7. For **Password Script**, leave the field blank.

8. For **Authentication type**, enter **password**. This field specifies the IPMI authentication type. Possible values for this field are none, **password**, **md2**, or **md5**.
9. Leave the **Use Lanplus** field blank.

After entering the configuration information for the backup fencing method for **clusternode2.example.com**, you can update the backup fence properties using the following procedure.

1. Click **Update backup fence properties** at the bottom of the right side of the screen. After the fence device has been added, a progress page is displayed after which the display returns to the configuration page for **clusternode2.example.com** in cluster **backupclust**.

To configure the fencing methods for **clusternode3.example.com**, use the same procedure as you did for configuring the fencing methods for **clusternode2.example.com**. In this case, however, use **3** as the port number for both of the APC switches that you are using for the main fencing method. For the backup fencing method, use **ipmifence3** as the name of the fence type and use an IP address of 10.15.86.52. The other components should be the same, as summarized in [Table 7.7, "Fence Agent Components to Specify for clusternode3.example.com"](#).

7.4. Cluster Configuration File for Backup Fence Method

Configuring a cluster with Conga modifies the cluster configuration file. This section shows the cluster configuration file before and after the procedures documented in [Section 7.3, "Backup Fencing Configuration Procedure"](#) were performed.

Before the cluster resources and service were configured, the **cluster.conf** file appeared as follows.

```
<?xml version="1.0"?>
<cluster alias="backupclust" config_version="34" name="backupclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
      <fence/>
    </clusternode>
  </clusternodes>
  <cman/>
  <fencedevices/>
  <rm>
    <failoverdomains/>
    <resources/>
  </rm>
</cluster>
```

After the cluster resources and service were configured, the `cluster.conf` file appears as follows.

```
<?xml version="1.0"?>
<cluster alias="backupclust" config_version="10" name="backupclust">
  <fence_daemon clean_start="0" post_fail_delay="0" post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1" votes="1">
      <fence>
        <method name="1">
          <device name="pwr01" option="off" port="1"/>
          <device name="pwr02" option="off" port="1"/>
          <device name="pwr01" option="on" port="1"/>
          <device name="pwr02" option="on" port="1"/>
        </method>
        <method name="2">
          <device name="ipmifence1"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2" votes="1">
      <fence>
        <method name="1">
          <device name="pwr01" option="off" port="2"/>
          <device name="pwr02" option="off" port="2"/>
          <device name="pwr01" option="on" port="2"/>
          <device name="pwr02" option="on" port="2"/>
        </method>
        <method name="2">
          <device name="ipmifence2"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="clusternode3.example.com" nodeid="3" votes="1">
      <fence>
        <method name="1">
          <device name="pwr01" option="off" port="3"/>
          <device name="pwr02" option="off" port="3"/>
          <device name="pwr01" option="on" port="3"/>
          <device name="pwr02" option="on" port="3"/>
        </method>
        <method name="2">
          <device name="ipmifence3"/>
        </method>
      </fence>
    </clusternode>
  </clusternodes>
  <cman/>
  <fencedevices>
    <fencedevice agent="fence_apc" ipaddr="10.15.86.96" login="apclogin"
name="pwr01" passwd="apcpword"/>
    <fencedevice agent="fence_apc" ipaddr="10.15.86.97" login="apclogin"
name="pwr02" passwd="apcpword"/>
    <fencedevice agent="fence_ipmilan" ipaddr="10.15.86.50"
login="ipmilogin" name="ipmifence1" passwd="ipmipword"/>
    <fencedevice agent="fence_ipmilan" ipaddr="10.15.86.51"
```

```
login="ipmilogin" name="ipmifence2" passwd="ipmipword"/>
    <fencedevice agent="fence_ipmilan" ipaddr="10.15.86.52"
login="ipmilogin" name="ipmifence3" passwd="ipmipword"/>
</fencedevices>
<rm>
    <failoverdomains/>
    <resources/>
</rm>
</cluster>
```

7.5. Testing the Backup Fence Device Configuration

To check whether the configuration you have defined works as expected, you can first try simply fencing the nodes to see if any issues arise:

```
# /sbin/fence_node clusternode1.example.com
# /sbin/fence_node clusternode2.example.com
# /sbin/fence_node clusternode3.example.com
```

To see whether the backup IPMI fencing will work when the primary APC switch fencing does not, disable the ethernet access to both APC switches. This will prevent the **fence_node** command from being able to access the switches. Then run the **fence_node** command on each node in the cluster to see whether the IPMI switch takes over and fences the node.

Chapter 8. Configuring Fencing using SCSI Persistent Reservations

This chapter provides the procedures for configuring fencing using SCSI persistent reservations in a Red Hat cluster using the Conga configuration tool.

SCSI persistent reservations provide the capability to control the access of each node to shared storage devices. Red Hat Cluster Suite employs SCSI persistent reservations as a fencing method through the use of the `fence_scsi` agent. The `fence_scsi` agent provides a method to revoke access to shared storage devices, provided that the storage support SCSI persistent reservations.

Using SCSI reservations as a fencing method is different from traditional power fencing methods. It is important to understand the software, hardware, and configuration requirements prior to using SCSI persistent reservations as a fencing method.

8.1. Technical Overview of SCSI Persistent Reservations

In order to understand how Red Hat Cluster Suite is able to use SCSI persistent reservations as a fencing method, it is helpful to have some basic knowledge of SCSI persistent reservations.

There are two important concepts withing SCSI persistent reservations that should be made clear: registrations and reservations.

8.1.1. SCSI Registrations

A registration occurs when a node registers a unique key with a device. A device can have many registrations. For our purposes, each node will create a registration on each device.

8.1.2. SCSI Technical Overview

A reservation dictates how a device can be accessed. In contrast to registrations, there can be only one reservation on a device at any time. The node that holds the reservation is know as the "reservation holder". The reservation defines how other nodes may access the device. For example, Red Hat Cluster Suite uses a "Write Exclusive, Registrants Only" reservation. This type of reservation indicates that only nodes that have registered with that device may write to the device.

8.1.3. SCSI Fencing with Persistent Reservations

Red Hat Cluster Suite is able to perform fencing via SCSI persistent reservations by simply removing a node's registration key from all devices. When a node failure occurs, the `fence_scsi` agent will remove the failed node's key from all devices, thus preventing it from being able to write to those devices.

8.2. SCSI Fencing Requirements and Limitations

In order to configure your system to use SCSI persistent reservations to fence a node, you must be sure that the following conditions are met.

- ✦ The `sg3_utils` package must be installed on your cluster nodes. This package provides the tools needed by the various scripts to manage SCSI persistent reservations.
- ✦ All shared storage must use LVM2 cluster volumes.
- ✦ All devices within the LVM2 cluster volumes must be SPC-3 compliant.

In addition to these requirements, fencing by way of SCSI persistent reservations is subject to the following limitations:

- All nodes in the cluster must have a consistent view of storage. Each node in the cluster must be able to remove another node's registration key from all the devices that it registered with. In order to do this, the node performing the fencing operation must be aware of all devices that other nodes are registered with.
- Devices used for the cluster volumes should be a complete LUN, not partitions. SCSI persistent reservations work on an entire LUN, meaning that access is controlled to each LUN, not individual partitions.
- As of Red Hat Enterprise Linux 5.5 and fully-updated releases of Red Hat Enterprise Linux 5.4, SCSI fencing can be used in a 2-node cluster; previous releases did not support this feature.
- As of Red Hat Enterprise Linux 5.5 and fully-updated releases of Red Hat Enterprise Linux 5.4, SCSI fencing can be used in conjunction with `qdisk`; previous releases did not support this feature. You cannot use `fence_scsi` on the LUN where `qdiskd` resides; it must be a raw LUN or raw partition of a LUN.

8.3. SCSI Fencing Example Configuration

[Figure 8.1, “Using SCSI Persistent Reservations as a Fence Device”](#) shows the configuration this procedure yields. All three nodes in this cluster have a consistent view of the storage, which means in this case that all of the nodes are registered with the same devices.

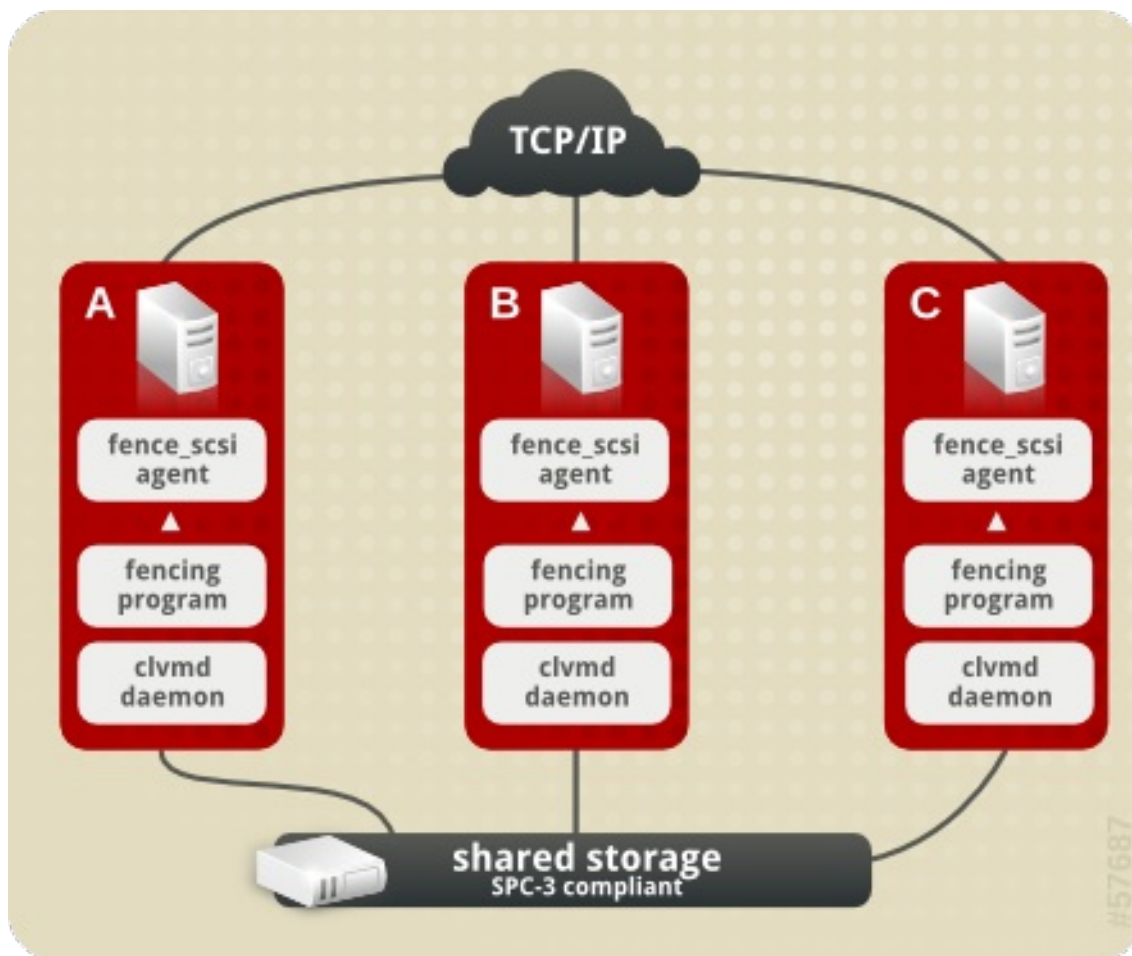


Figure 8.1. Using SCSI Persistent Reservations as a Fence Device

8.4. SCSI Fencing Prerequisite Configuration

[Table 8.1, “Configuration Prerequisites”](#) summarizes the prerequisite components that have been set up before this procedure begins.

Table 8.1. Configuration Prerequisites

Component	Name	Comment
cluster	scsiclust	three-node cluster
cluster node	clusternode1.example.com	node in cluster scsiclust with sg3_utils package installed
cluster node	clusternode2.example.com	node in cluster scsiclust with sg3_utils package installed
cluster node	clusternode3.example.com	node in cluster scsiclust with sg3_utils package installed

8.5. SCSI Fence Device Components to Configure

This procedure configures SCSI persistent reservations as a fence method for each node in cluster **scsiclust**.

[Table 8.2, “Fence Agent Components to Configure for clusternode1.example.com”](#) summarizes the components of the SCSI fence device that this procedure configures for cluster node **clusternode1.example.com**.

Table 8.2. Fence Agent Components to Configure for clusternode1.example.com

Fence Agent Component	Value	Description
Name	scsifence	name of the SCSI fencing device
Node name	node1	name of node to be fenced

[Table 8.3, “Fence Agent Components to Configure for clusternode2.example.com”](#) summarizes the components of the scsi fence device that this procedure configures for cluster node **clusternode2.example.com**.

Table 8.3. Fence Agent Components to Configure for clusternode2.example.com

Fence Agent Component	Value	Description
Name	scsifence	name of the SCSI fencing device
Node name	node2	name of node to be fenced

[Table 8.4, “Fence Agent Components to Configure for clusternode3.example.com”](#) summarizes the components of the SCSI fence device that this procedure configures for cluster node **clusternode3.example.com**.

Table 8.4. Fence Agent Components to Configure for clusternode3.example.com

Fence Agent Component	Value	Description
Name	scsifence	name of the SCSI fencing device
Node name	node3	name of node to be fenced

8.6. SCSI Fence Device Configuration Procedure

This section provides the procedure for configuring SCSI persistent reservations as a fencing mechanism for each node of cluster **scsiclust**.

Use the following procedure to configure the HP iLO management board as the fence device for node **clusternode1.example.com** using Conga:

1. As an administrator of **luci** Select the **cluster** tab. This displays the **Choose a cluster to administer** screen.
2. From the **Choose a cluster to administer** screen, you should see the previously configured cluster **scsiclust** displayed, along with the nodes that make up the cluster. Click on **clusternode1.example.com**. This displays the configuration screen for node **clusternode1.example.com**.
3. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
4. From the dropdown menu, under **Create a new Fence Device**, select **SCS fencing**. This displays a fence device configuration menu.
5. For **Name**, enter **scsifence**.
6. For **Node name**, enter **node1**.
7. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
8. On the confirmation screen, click **OK**. After the fence device has been added, a progress page is displayed after which the display returns to the configuration page for **clusternode1.example.com** in cluster **scsiclust**.

After configuring a SCSI fence device for **clusternode1.example.com**, use the following procedure to configure a SCSI fence device for **clusternode2.example.com**.

1. From the configuration page for **clusternode1.example.com**, a menu appears on the left of the screen for cluster **scsiclust**. Select the node **clusternode2.example.com**. The configuration page for **clusternode2.example.com** appears, with no fence device configured.
2. At the **Main Fencing Method** display, click **Add a fence device to this level**. This causes a dropdown menu to display.
3. From the dropdown menu, under **Use an existing Fence Device**, you should see **scsifence (SCSI Reservation)**, which you defined for **clusternode1.example.com**. Select this existing device, which displays a fence device configuration menu.
4. For **Node Name**, enter **node2**.
5. Click **Update main fence properties**. This causes a confirmation screen to be displayed.
6. On the confirmation screen, click **OK**. After the fence device has been added, a progress page is displayed after which the display returns to the configuration page for **clusternode2.example.com** in cluster **scsiclust**.

After configuring **scsifence** as the fencing device for **clusternode2.example.com**, select node **clusternode3.example.com** from the menu on the left side of the page and configure a SCSI fence device for that node using the same procedure as you did to configure the fence devices for **clusternode2.example.com**. For **clusternode3.example.com**, use the existing fence method

scsifence as the name of the fencing method and **node3** as the host name.

8.7. Cluster Configuration File with SCSI Fence Device

Configuring a cluster with Conga modifies the cluster configuration file. This section shows the cluster configuration file before and after the procedures documented in [Section 8.6, “SCSI Fence Device Configuration Procedure”](#) and were performed.

Before the cluster resources and service were configured, the **cluster.conf** file appeared as follows.

```
<?xml version="1.0"?>
<cluster alias="scsiclust" config_version="12" name="scsiclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
      <fence/>
    </clusternode>
    <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
      <fence>
        <method name="1"/>
      </fence>
    </clusternode>
  </clusternodes>
  <cman/>
  <fencedevices/>
  <rm>
    <failoverdomains/>
    <resources/>
  </rm>
</cluster>
```

After the cluster resources and service were configured, the **cluster.conf** file appears as follows.

```
<?xml version="1.0"?>
<cluster alias="scsiclust" config_version="19" name="scsiclust">
  <fence_daemon clean_start="0" post_fail_delay="0"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="clusternode1.example.com" nodeid="1"
votes="1">
      <fence>
        <method name="1">
          <device name="scsifence"
node="node1"/>
        </method>
```

```

        </fence>
    </clusternode>
    <clusternode name="clusternode2.example.com" nodeid="2"
votes="1">
        <fence>
            <method name="1">
                <device name="scsifence"
node="node2"/>
            </method>
        </fence>
    </clusternode>
    <clusternode name="clusternode3.example.com" nodeid="3"
votes="1">
        <fence>
            <method name="1">
                <device name="scsifence"
node="node3"/>
            </method>
        </fence>
    </clusternode>
</clusternodes>
<cman/>
<fencedevices>
    <fencedevice agent="fence_scsi" name="scsifence"/>
</fencedevices>
<rm>
    <failoverdomains/>
    <resources/>
</rm>
</cluster>

```

8.8. Testing the Configuration

It is important that all SCSI fencing requirements be met in order for your system to successfully fence a node using SCSI persistent reservations. The SCSI fencing requirements are noted in [Section 8.2, “SCSI Fencing Requirements and Limitations”](#). To ensure that your system meets these requirements, you should test your configuration.

After the **cluster.conf** has been set up on all of the nodes in the system, you can perform the following procedure to verify that all of the requirements have been met for SCSI fencing and that the configuration is successful.

1. For every node in the cluster, you should verify that the necessary infrastructure is up and running:
 - ✦ Ensure that the cluster infrastructure is up and running on every node in the cluster; you can check this with the **cman_tool status** command.
 - ✦ Ensure that the **clvmd** daemon is running; you can check this with the **service clvmd status** command.
 - ✦ Ensure that the **scsi_reserve** service has been turned on by executing the **chkconfig scsi_reserve on** command.
2. Set up cluster LVM volumes to test.

```
[root@tng3-1 ~]# pvcreate /dev/sda1 /dev/sdb1 /dev/sdc1
```

```
[root@tng3-1 ~]# vgcreate new_vol_group /dev/sda1 /dev/sdb1 /dev/sdc1
```

```
[root@tng3-1 ~]# lvcreate -L2G -n new_logical_volume new_vol_group
```

```
[root@tng3-1 ~]# gfs_mkfs -plock_nolock -j 1  
/dev/new_vol_group/new_logical_volume
```

```
[root@tng3-1 ~]# mount /dev/new_vol_group/new_logical_volume /mnt
```

3. Run the **scsi_reserve init** script on all nodes, and then check to see whether this worked.

```
[root@clusternode1 ~]# service scsi_reserve start  
[root@clusternode1 ~]# service scsi_reserve status  
[root@clusternode2 ~]# service scsi_reserve start  
[root@clusternode2 ~]# service scsi_reserve status  
[root@clusternode3 ~]# service scsi_reserve start  
[root@clusternode3 ~]# service scsi_reserve status
```

4. Execute the following commands and check whether the nodes have been fenced.

```
# /sbin/fence_node clusternode1.example.com  
# /sbin/fence_node clusternode2.example.com  
# /sbin/fence_node clusternode3.example.com
```

Chapter 9. Troubleshooting

The following is a list of some problems you may see regarding the configuration of fence devices as well as some suggestions for how to address these problems.

- If your system does not fence a node automatically, you can try to fence the node from the command line using the **fence_node** command, as described at the end of each of the fencing configuration procedures. The **fence_node** performs I/O fencing on a single node by reading the fencing settings from the **cluster.conf** file for the given node and then running the configured fencing agent against the node. For example, the following command fences node **clusternode1.example.com**:

```
# /sbin/fence_node clusternode1.example.com
```

If the **fence_node** command is unsuccessful, you may have made an error in defining the fence device configuration. To determine whether the fencing agent itself is able to talk to the fencing device, you can execute the I/O fencing command for your fence device directly from the command line. As a first step, you can execute the with the **-o status** option specified. For example, if you are using an APC switch as a fencing agent, you can execute a command such as the following:

```
# /sbin/fence_apc -a (ipaddress) -l (login) ... -o status -v
```

You can also use the I/O fencing command for your device to fence the node. For example, for an HP ILO device, you can issue the following command:

```
# /sbin/fence_ilo -a myilo -l login -p passwd -o off -v
```

- Check the version of firmware you are using in your fence device. You may want to consider upgrading your firmware. You may also want to scan bugzilla to see if there are any issues regarding your level of firmware.
- If a node in your cluster is repeatedly getting fenced, it means that one of the nodes in your cluster is not seeing enough "heartbeat" network messages from the node that is getting fenced. Most of the time, this is a result of flaky or faulty hardware, such as bad cables or bad ports on the network hub or switch. Test your communications paths thoroughly without the cluster software running to make sure your hardware is working correctly.
- If a node in your cluster is repeatedly getting fenced right at startup, it may be due to system activities that occur when a node joins a cluster. If your network is busy, your cluster may decide it is not getting enough heartbeat packets. To address this, you may have to increase the **post_join_delay** setting in your **cluster.conf** file. This delay is basically a grace period to give the node more time to join the cluster.

In the following example, the **fence_daemon** entry in the cluster configuration file shows a **post_join_delay** setting that has been increased to 600.

```
<fence_daemon clean_start="0" post_fail_delay="0" post_join_delay="600">
```

- If a node fails while the **fenced** daemon is not running, it will not be fenced. It will cause problems if the **fenced** daemon is killed or exits while the node is using GFS. If the **fenced** daemon exits, it should be restarted.

If you find that you are seeing error messages when you try to configure your system, or if after configuration your system does not behave as expected, you can perform the following checks and examine the following areas.

- ✦ Connect to one of the nodes in the cluster and execute the **clustat(8)** command. This command runs a utility that displays the status of the cluster. It shows membership information, quorum view, and the state of all configured user services.

The following example shows the output of the **clustat(8)** command.

```
[root@clusternode4 ~]# clustat
Cluster Status for nfsclust @ Wed Dec  3 12:37:22 2008
Member Status: Quorate

Member Name                ID  Status
-----
clusternode5.example.com   1  Online, rgmanager
clusternode4.example.com   2  Online, Local, rgmanager
clusternode3.example.com   3  Online, rgmanager
clusternode2.example.com   4  Online, rgmanager
clusternode1.example.com   5  Online, rgmanager

Service Name                Owner (Last)                State
-----
service:nfssvc              clusternode2.example.com    starting
```

In this example, **clusternode4** is the local node since it is the host from which the command was run. If **rgmanager** did not appear in the **Status** category, it could indicate that cluster services are not running on the node.

- ✦ Connect to one of the nodes in the cluster and execute the **group_tool(8)** command. This command provides information that you may find helpful in debugging your system. The following example shows the output of the **group_tool(8)** command.

```
[root@clusternode1 ~]# group_tool
type          level name      id          state
fence         0    default  00010005   none
[1 2 3 4 5]
dlm           1    clvmd    00020005   none
[1 2 3 4 5]
dlm           1    rgmanager 00030005   none
[3 4 5]
dlm           1    mygfs    007f0005   none
[5]
gfs           2    mygfs    007e0005   none
[5]
```

The state of the group should be **none**. The numbers in the brackets are the node ID numbers of the cluster nodes in the group. The **clustat** shows which node IDs are associated with which nodes. If you do not see a node number in the group, it is not a member of that group. For example, if a node ID is not in dlm/rgmanager group, it is not using the rgmanager dlm lock space (and probably is not running rgmanager).

The level of a group indicates the recovery ordering. 0 is recovered first, 1 is recovered second, and so forth.

- ✦ Connect to one of the nodes in the cluster and execute the **cman_tool nodes -f** command. This command provides information about the cluster nodes that you may want to look at. The following example shows the output of the **cman_tool nodes -f** command.

```
[root@clusternode1 ~]# cman_tool nodes -f
Node  Sts   Inc   Joined                               Name
  1    M    752   2008-10-27 11:17:15 clusternode5.example.com
  2    M    752   2008-10-27 11:17:15 clusternode4.example.com
  3    M    760   2008-12-03 11:28:44 clusternode3.example.com
  4    M    756   2008-12-03 11:28:26 clusternode2.example.com
  5    M    744   2008-10-27 11:17:15 clusternode1.example.com
```

The **Sts** heading indicates the status of a node. A status of M indicates the node is a member of the cluster. A status of X indicates that the node is dead. The **Inc** heading indicating the incarnation number of a node, which is for debugging purposes only.

- » Check whether the **cluster.conf** is identical in each node of the cluster. If you configure your system with Conga, as in the example provided in this document, these files should be identical, but one of the files may have accidentally been deleted or altered.

Chapter 10. The GFS Withdraw Function

The GFS *withdraw* function is a data integrity feature of GFS file systems in a cluster. If the GFS kernel module detects an inconsistency in a GFS file system following an I/O operation, the file system becomes unavailable to the cluster. The I/O operation stops and the system waits for further I/O operations to error out, preventing further damage. When this occurs, you can stop any other services or applications manually, after which you can reboot and remount the GFS file system to replay the journals. If the problem persists, you can unmount the file system from all nodes in the cluster and perform file system recovery with the **gfs_fsck** command. The GFS withdraw function is less severe than a kernel panic, which would cause another node to fence the node.

You can override the GFS withdraw function by mounting the file system with the **-o errors=panic** option specified. When this option is specified, any errors that would normally cause the system to withdraw cause the system to panic instead. This stops the node's cluster communications, which causes the node to be fenced.

For information on the GFS withdraw function, see *Global File System: Configuration and Administration*.

Appendix A. Revision History

Revision 2-19.33.400	2013-10-31	Rüdiger Landmann
Rebuild with publican 4.0.0		
Revision 2-19.33	July 24 2012	Ruediger Landmann
Rebuild for Publican 3.0		
Revision 5.6-1	Thu Dec 10 2010	Steven Levine
Resolves: #581594 Updates information about support for qdisk with SCSI reservations.		
Resolves: #622550 Updates information about support for SCSI reservations in a 2-node cluster.		
Revision 2.0-0	Mon Mar 15 2010	Steven Levine
Added configuration examples for additional scenarios		
Revision 1.0-0	Thu Jun 17 2009	Steven Levine
First edition		

Index

A

APC fence device configuration

- components to configure, [APC Fence Device Components to Configure](#), [Fence Device Components to Configure](#)
- prerequisites, [APC Fence Device Prerequisite Configuration](#)
- procedure, [APC Fence Device Configuration Procedure](#)

APC switch

- configuring as fence device, [Configuring an APC Switch as a Fence Device](#)
- configuring as sharable fence device, [APC Fence Device Configuration Procedure](#)
- testing fence configuration, [Testing the APC Fence Device Configuration](#)

APC switch configuration component

- IP Address, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- Login, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- Name, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- Password, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- Password Script, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- Port, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)
- Switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)
- Use SSH, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#)

[Procedure, Configuring Fencing on the First Cluster Node](#)**Authentication Type configuration component**

- HP iLO board, [HP iLO Fence Device Configuration Procedure](#)
- IPMI board, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

B**backup fence configuration**

- prerequisites, [Backup Fencing Prerequisite Configuration](#)

backup fence method

- testing fence configuration, [Testing the Backup Fence Device Configuration](#)

backup fencing configuration, [Configuring Fencing on the First Cluster Node](#)

- procedure, [Backup Fencing Configuration Procedure](#)

Backup Fencing Method configuration, [Configuring Fencing on the First Cluster Node](#)**C****clustat command, [Troubleshooting](#)**

cluster.conf file, [Cluster Configuration File with APC Fence Device](#), [Cluster Configuration File with IPMI Fence Device](#), [Cluster Configuration File with HP iLO Fence Device](#), [Cluster Configuration File with Dual Power Supply Fencing](#), [Cluster Configuration File for Backup Fence Method](#), [Cluster Configuration File with SCSI Fence Device](#)

cman_tool command, [Troubleshooting](#)**D****dual power**

- testing fence configuration, [Testing the Dual Power Fence Device Configuration](#)

dual power fence configuration

- prerequisites, [Dual Power Fencing Prerequisite Configuration](#)

dual power fencing configuration, [Dual Power Fencing Configuration Procedure](#)

- components to configure, [Fence Device Components to Configure](#), [Fence Device Components to Configure](#)
- procedure, [Dual Power Fencing Configuration Procedure](#)

F**fence device**

- APC switch, [Configuring an APC Switch as a Fence Device](#)
- backup, [Configuring a Backup Fencing Method](#)
- dual power, [Configuring Fencing with Dual Power Supplies](#)
- HP iLO management board, [Configuring HP iLO Management Boards as Fencing Devices](#)
- IPMI management board, [Configuring IPMI Management Boards as Fencing Devices](#)
- SCSI persistent reservations, [Configuring Fencing using SCSI Persistent Reservations](#)

fence_apc command, [Troubleshooting](#)**fence_ilo command, [Troubleshooting](#)**

fence_node command, [Testing the APC Fence Device Configuration](#), [Testing the IPMI Fence Device Configuration](#), [Testing the HP iLO Fence Device Configuration](#), [Testing the Dual Power Fence Device Configuration](#), [Testing the Configuration](#), [Troubleshooting](#)

G

GFS withdraw function, [The GFS Withdraw Function](#)

group_tool command, [Troubleshooting](#)

H**HP iLO board configuration component**

- Authentication Type, [HP iLO Fence Device Configuration Procedure](#)
- IP Address, [HP iLO Fence Device Configuration Procedure](#)
- Login, [HP iLO Fence Device Configuration Procedure](#)
- Name, [HP iLO Fence Device Configuration Procedure](#)
- Password, [HP iLO Fence Device Configuration Procedure](#)
- Password Script, [HP iLO Fence Device Configuration Procedure](#)
- Use Lanplus, [HP iLO Fence Device Configuration Procedure](#)

HP iLO fence device configuration

- components to configure, [HP iLO Fence Device Components to Configure](#)
- prerequisites, [HP iLO Fence Device Prerequisite Configuration](#)
- procedure, [HP iLO Fence Device Configuration Procedure](#)

HP iLO management board

- configuring as fence device, [Configuring HP iLO Management Boards as Fencing Devices](#)
- testing fence configuration, [Testing the HP iLO Fence Device Configuration](#)

I**IP Address configuration component**

- APC switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- HP iLO board, [HP iLO Fence Device Configuration Procedure](#)
- IPMI board, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

IPMI board configuration component

- Authentication Type, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)
- IP Address, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)
- Login, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)
- Name, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)
- Password, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)
- Password Script, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)
- Use Lanplus, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

IPMI fence device configuration

- components to configure, [IPMI Fence Device Components to Configure](#), [Fence Device Components to Configure](#)
- prerequisites, [IPMI Fence Device Prerequisite Configuration](#)
- procedure, [IPMI Fence Device Configuration Procedure](#)

IPMI management board

- configuring as fence device, [Configuring IPMI Management Boards as Fencing Devices](#)
- testing fence configuration, [Testing the IPMI Fence Device Configuration](#)

L

Login configuration component

- APC switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- HP iLO board, [HP iLO Fence Device Configuration Procedure](#)
- IPMI board, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

M

Main Fencing Method configuration, [APC Fence Device Configuration Procedure](#), [IPMI Fence Device Configuration Procedure](#), [HP iLO Fence Device Configuration Procedure](#), [SCSI Fence Device Configuration Procedure](#)**main fencing method configuration, [Configuring Fencing on the First Cluster Node](#), [Configuring Fencing on the Remaining Cluster Nodes](#)**

N

Name configuration component

- APC switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- HP iLO board, [HP iLO Fence Device Configuration Procedure](#)
- IPMI board, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

P

Password configuration component

- APC switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- HP iLO board, [HP iLO Fence Device Configuration Procedure](#)
- IPMI board, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

Password Script configuration component

- APC switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring the APC switches as shared fence devices](#)
- HP iLO board, [HP iLO Fence Device Configuration Procedure](#)
- IPMI board, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

Port configuration component

- APC switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

post_join_delay setting in cluster.conf, [Troubleshooting](#)

S

SCSI fence device configuration

- components to configure, [SCSI Fence Device Components to Configure](#)
- prerequisites, [SCSI Fencing Prerequisite Configuration](#)
- procedure, [SCSI Fence Device Configuration Procedure](#)

SCSI persistent reservations

- configuring as fence device, [Configuring Fencing using SCSI Persistent Reservations](#)

sharable fence device

- configuration, [APC Fence Device Configuration Procedure](#)

Switch configuration component

- APC switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

T

testing fence configuration

- APC switch, [Testing the APC Fence Device Configuration](#)
- backup method, [Testing the Backup Fence Device Configuration](#)
- dual power, [Testing the Dual Power Fence Device Configuration](#)
- HP iLO management board, [Testing the HP iLO Fence Device Configuration](#)
- IPMI management board, [Testing the IPMI Fence Device Configuration](#)

U

Use Lanplus configuration component

- HP iLO board, [HP iLO Fence Device Configuration Procedure](#)
- IPMI board, [IPMI Fence Device Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

Use SSH configuration component

- APC switch, [APC Fence Device Configuration Procedure](#), [Dual Power Fencing Configuration Procedure](#), [Configuring Fencing on the First Cluster Node](#)

W

withdraw function, GFS, [The GFS Withdraw Function](#)