



# **Multi-System High Availability Configuration Guide**

**Version #: XOS 9.0.0**

---

## Copyright and Trademark Information

Copyright © 2010 by Crossbeam Systems®

Boxborough, MA, USA

All Rights Reserved

The products, specifications, and other technical information regarding the products contained in this document are subject to change without notice. All information in this document is believed to be accurate and reliable, but is presented without warranty of any kind, expressed or implied, and users must take full responsibility for their application of any products specified in this document. Crossbeam Systems disclaims responsibility for errors that may appear in this document, and it reserves the right, in its sole discretion and without notice, to make substitutions and modifications in the products and practices described in this document.

This material is protected by the copyright and trade secret laws of the United States and other countries. It may not be reproduced, distributed, or altered in any fashion by any entity (either internal or external to Crossbeam Systems), except in accordance with applicable agreements, contracts, or licensing, without the express written consent of Crossbeam Systems.

For permission to reproduce or distribute please contact your Crossbeam Systems account executive.

This product includes software developed by the Apache Software Foundation: <http://www.apache.org>.

“Crossbeam,” “Crossbeam Systems,” “iBeam,” “XOS,” X40, X45, X80, “COS,” C2, C6, C12, C25, and any logos associated therewith are trademarks or registered trademarks of Crossbeam Systems, Inc. in the U.S. Patent and Trademark Office, and several international jurisdictions.

All other product names mentioned in this manual may be trademarks or registered trademarks of their respective companies.

---

---

# Contents

## About This Guide

Intended Audience .....	5
Related Documentation .....	5
Software Documentation .....	5
Conventions .....	6
Typographical Conventions .....	6
Cautions, Warnings, and Notes .....	7
Customer Support .....	8

## Chapter 1: Configuring for High Availability

Overview of High Availability Mechanisms .....	9
High Availability Within an X-Series Chassis .....	9
CPM Redundancy .....	9
NPM Interface Redundancy .....	9
APM Redundancy .....	9
Disk Redundancy .....	9
High Availability on Multiple X-Series Chassis .....	9
Virtual Router Redundancy Protocol (VRRP) .....	9
Open Shortest Path First (OSPF) Cost .....	10
Active-Active Configuration .....	10
Active-Standby Configuration .....	10
Failover Groups .....	10
Virtual Router (VR) .....	10
VRRP Priority .....	10
Control Link Port (HA Port) .....	11
Configuring an Active-Active High Availability System .....	11
Configuring an Active-Standby High Availability System .....	14

## Chapter 2: Active-Standby VRRP Dual-box, High Availability Configuration

Chassis Hardware Configurations .....	17
Assumptions .....	18
Active-Standby Configuration .....	18
System Diagram (Active-Standby) .....	19
Configuring the virtual-router for the gig_16 circuit .....	22
Configuring the virtual-router for the gig_26 circuit .....	23
Verifying Your Configuration .....	25
Output of show running-config on Chassis 1 .....	25
Output of show vrrp on Chassis 1 .....	26
Output of show vrrp detail-status on Chassis 1 .....	26
Output of show remote-box on Chassis 1 .....	26
Configuring the virtual-router for the gig_14 circuit .....	29
Configuring the virtual-router for the gig_24 circuit .....	30
Verifying Your Configuration .....	32
Output of show running-config on Chassis 2 .....	32
Output of show vrrp on Chassis 2 .....	33
Output of show vrrp detail-status on Chassis 2 .....	33
Output of show remote-box on Chassis 2 .....	33

## Chapter 3: Active-Active VRRP Dual-box, High Availability Configuration

Chassis Hardware Configurations .....	35
---------------------------------------	----

Assumptions .....	36
Active-Active Configuration .....	36
System Diagram (Active-Active) .....	37
Configuring the virtual-router for the gig_18 circuit .....	40
Configuring the virtual-router for the gig_28 circuit .....	41
Configuring the virtual-router for the gig_18 circuit .....	44
Configuring the virtual-router for the gig_28 circuit .....	45
Verifying Your Configuration .....	47
Output of show running-config on Chassis 1 .....	47
Output of show vrrp on Chassis 1 .....	48
Output of show vrrp detail-status on Chassis 1 .....	49
Output of show remote-box on Chassis 1 .....	49
Configuring the virtual-router for the gig_17 circuit .....	52
Configuring the virtual-router for the gig_27 circuit .....	53
Configuring the virtual-router for the gig_17 circuit .....	56
Configuring the virtual-router for the gig_27 circuit .....	56
Verifying Your Configuration .....	59
Output of show running-config on Chassis 2 .....	59
Output of show vrrp on Chassis 2 .....	60
Output of show vrrp detail-status on Chassis 2 .....	61
Output of show remote-box on Chassis 1 .....	61

## Appendix A: Basic Chassis Configuration

Assign Hostnames .....	63
Chassis 1 .....	63
Chassis 2 .....	63
Assign a Domain Name .....	63
Chassis 1 .....	63
Chassis 2 .....	63

---

# About This Guide

This guide provides step-by-step instructions for creating high availability X-Series Platform configurations running XOS Version 9.0 or later.

This guide assumes that you have already installed the X-Series Platform hardware, and that you have a basic understanding of how the X-Series Platform is designed and operates.

## Intended Audience

This guide is intended for system integrators and other qualified service personnel responsible for installing, configuring, and managing the Crossbeam X-Series Platform.

**IMPORTANT:** For the latest updates and revisions to X-Series Platform documentation, log into the Crossbeam Online Support Portal at <http://www.crossbeam.com/support/online-support/>.

## Related Documentation

The following documents are provided on the Crossbeam Systems Documentation DVD and the Crossbeam Systems Customer Support Web site at <http://www.crossbeam.com/support/online-support/>.

### Software Documentation

- *XOS Configuration Guide*
- *XOS Command Reference Guide*
- *XOS Release Notes*
- *X80 Platform Hardware Installation Guide*
- *X45 Platform Hardware Installation Guide*
- *NPM-8600 Installation Notice*
- *Install Server User Guide*
- *Install Server Release Notes*
- *RSW Installation Guide*
- *RSW Release Notes*
- *Check Point® Security Gateway R70 and Check Point® Security Gateway R71 Installation and Configuration Guide for Crossbeam X-Series Platforms*
- *Check Point® VPN-1 Power NGX R65 Installation and Configuration Guide for Crossbeam® X-Series Platforms*
- *Check Point® VPN-1 Power VSX NGX R65 Installation and Configuration Guide for Crossbeam® X-Series*
- *Check Point® VSX NGX R67 Installation and Configuration Guide for Crossbeam® X-Series Platforms*

# Conventions

## Typographical Conventions

For paragraph text conventions, see [Typographical Conventions Used in Paragraph Text](#) on page 6

For command-line text conventions, see [Typographical Conventions Used in Command-Line Text](#) on page 7

**Table 1. Typographical Conventions Used in Paragraph Text**

Typographical Convention	Types of Information	Usage Examples
<b>Bold</b>	Elements on the graphical user interface.	In the <b>IP Address</b> field, type the IP address of the first VAP in the group. Click <b>OK</b> to close the dialog. Select the <b>Print to File</b> check box.
<i>Courier</i>	Keys on the keyboard. File names, folder names, and command names. Any information that you must type exactly as shown. Program output text.	Press <code>Esc</code> to return to the main menu. Save the <code>user.txt</code> file in the <code>user_install</code> directory. Use the <code>start</code> command to start the application. In the <b>Username</b> field, type <code>Administrator</code> . The XOS CLI <code>show calendar</code> command displays the system calendar: <code>Fri Mar 7 13:32:03 2009</code>
<i>Courier Italic</i>	File names, folder names, command names, or other information that you must supply.	In the <b>Version Number</b> field, type <code>8.5.patch_number</code> .
>	A sequence of commands from the task bar or menu bar.	From the taskbar, choose <b>Start &gt; Run</b> . From the main menu, choose <b>File &gt; Save As...</b> Right-click on the desktop and choose <b>Arrange Icons By &gt; Name</b> from the pop-up menu.

**Table 2. Typographical Conventions Used in Command-Line Text**

Typographical Convention	Types of Information	Usage Examples
Courier	User prompts and program output text.	CBS# <b>show calendar</b> Fri Mar 7 13:32:03 2008
<b>Courier Bold</b>	Information that you must type in exactly as shown.	[root@xxxxxx]# <b>md crossbeam</b>
< <i>Courier Italic</i> >	Angle brackets surrounding Courier italic text indicate file names, folder names, command names, or other information that you must supply.	[root@xxxxxx]# <b>md</b> <your_folder_name>
[ ]	Square brackets contain optional information that may be supplied with a command.	[root@xxxxxx]# <b>dir</b> [drive:] [path] [<filename>] [/P] [/W] [/D].
	Separates two or more mutually exclusive options.	[root@xxxxxx]# <b>verify</b> [ON OFF]
{ }	Braces contain two or more mutually exclusive options from which you must choose one.	CBS# <b>configure vap-group</b> <VAP_group_name> CBS(config-vap-grp)# <b>raid</b> {0 1}

## Cautions, Warnings, and Notes



**Caution:** Lists precautions that you must take to avoid temporary data loss or data unavailability.



**Warning:** Lists precautions that you must take to avoid personal injury, permanent data loss, or equipment damage.



**Electrical Hazard:** Lists precautions that you must take to avoid an electrical hazard that may result in personal injury, up to and including loss of life.

**IMPORTANT:** Lists important steps that you must perform properly or important information that you must take into consideration to avoid performing unnecessary work.

**NOTE:** Provides special information or tips that help you properly understand or carry out a task.

# Customer Support

Crossbeam Systems offers a variety of service plans designed to meet your specific technical support requirements. For information on purchasing a service plan for your organization, please contact your account representative or refer to <http://www.crossbeam.com/support/technical-support/>.

If you have purchased a Crossbeam Systems product service plan and need technical assistance, you can report issues by telephone:

**United States:** +1 800-331-1338 **OR** +1 978-318-7595

**EMEA:** + 33 4 8986 0400 (during normal working hours)

+1 978-318-7595 (outside office hours and on public holidays, if applicable)

**Asia Pacific:** +1 978-318-7595

**Latin America:** +1 978-318-7595

You can also report issues via e-mail to [support@crossbeamsystems.com](mailto:support@crossbeamsystems.com).

In addition, all of our service plans include access to the Crossbeam online support Web site located at <http://www.crossbeam.com/support/online-support/>.

The Crossbeam online support Web site provides you with access to a variety of resources, including Customer Support Knowledgebase articles, technical bulletins, product documentation, and release notes. You can also access our real-time problem reporting application, which lets you submit new technical support requests and view all your open requests.

Crossbeam Systems also offers extensive customer training on all of its products. For current course offerings and schedules, please refer to the Crossbeam training and education Web site located at <http://www.crossbeam.com/support/training-services/>.

---

# Configuring for High Availability

This chapter provides detailed information about setting up X-Series platforms to achieve the most common High Availability configurations. The following topics are covered in detail:

- [Overview of High Availability Mechanisms](#) on page 9
- [Configuring an Active-Active High Availability System](#) on page 11
- [Configuring an Active-Standby High Availability System](#) on page 14

## Overview of High Availability Mechanisms

High Availability is implemented in several ways on X-Series chassis.

### High Availability Within an X-Series Chassis

#### CPM Redundancy

Within an X-Series chassis, two CPMs can be configured to operate as a primary-backup pair.

#### NPM Interface Redundancy

On an NPM, an interface can be defined as a backup to one or more master interfaces.

#### APM Redundancy

APMs can be configured in standby mode, ready to substitute for an APM that fails.

VAP groups that experience a VAP failure, can be configured to preemptively acquire an APM from a lower-priority VAP group.

#### Disk Redundancy

On CPMs and APMs, pairs of disks can be configured in RAID 0 or RAID 1 arrays.

### High Availability on Multiple X-Series Chassis

#### Virtual Router Redundancy Protocol (VRRP)

Using VRRP, two or more X-Series chassis can be configured so that network traffic is re-routed from a primary chassis to a standby chassis if a failure or unwanted change occurs on the primary chassis in any of these areas:

- Circuit
- Interface (physical or logical)
- VAP Group

- IP Address
- Next Hop IP Address
- Multi-Link Trunking Interface

## Open Shortest Path First (OSPF) Cost

Another aspect of a failover involves the adjustment of certain parameters when a failover condition occurs. For example, the OSPF link cost associated with a failover group circuit can be increased when that group changes from master to backup. The Crossbeam Routing Software (RSW) then updates OSPF routes to ensure that traffic is routed through a circuit that is associated with the failover group that has become master.

## Active-Active Configuration

In an Active-Active configuration, both X-Series chassis handle traffic. If a failure occurs on either one, the remaining chassis takes over the traffic processing for both. When the failure condition is resolved, the traffic routing reverts to what it was before the failure. Both chassis must have sufficient processing capacity to handle the total workload.

## Active-Standby Configuration

In an Active-Standby configuration, all of the traffic is handled by the primary X-Series chassis. The standby chassis handles no traffic until a failover occurs, at which time the standby chassis assumes the primary role. Later, when the failure condition is resolved, the original roles may or may not be restored, depending on the configuration.

## Failover Groups

Failover groups and Virtual Routers (VRs) are used only in High Availability configurations. A failover group is a grouping of one or more VRs. A VR identifies the circuits and associated VAP groups for high availability. Only a failover group, not the entire system or an individual VAP group, can fail over to a standby failover group on another system.

## Virtual Router (VR)

A virtual router can be attached to a single circuit only, and can include only one VAP group attached to that circuit. In addition, the VR can assign individual IP addresses to the circuit and the VAP group interface. For circuits already configured with an IP address, the VR can also assign a virtual IP address. This virtual IP address allows you to configure failover groups using the same virtual IP address on other systems. Each virtual router can be configured to verify the state of the next hop IP address before using it.

## VRRP Priority

Each failover group is assigned a VRRP priority. Typically failover groups are defined in pairs and the failover group with the higher priority is designated the Master. Both failover groups in a pair must have the same ID and must be configured with different VRRP priorities.

A failure within a chassis does not necessarily cause a failover from one failover group to another. Instead, the VRRP priority is reduced by a pre-configured value, called a priority-delta. Failover occurs only if the priority is reduced below the priority of the backup failover group. This minimizes or eliminates the problem of failing over to a chassis that has an even more diminished capacity. After any failure is rectified, the VRRP priority is increased by the same amount by which it was decremented when that failure occurred. When all failures are rectified, the priority returns to the originally configured value.

## Control Link Port (HA Port)


The X-Series Platform requires a communication link between all the X-Series Platforms in a High Availability (HA) configuration. This is provided by the Control Link port on the CPM. In a dual-system configuration, the Control Link ports are connected to the same network broadcast domain. Crossbeam recommends that customers do not connect the Control Link ports directly to each other. For 3 or more systems, the Control Link ports are connected to a switch.

**NOTE:** Typically, communication between the control link ports is configured to use auto-negotiation. If you use a switch to connect the High Availability (HA) ports and auto-negotiation does not work, use the **configure management high-availability** command to manually set up the communication parameters.

## Configuring an Active-Active High Availability System

In an Active-Active configuration, each system processes traffic and each has two failover groups configured. On each system, you attach one circuit to two Virtual Routers (VRs), where each VR is in a different failover group. The basic configuration, is shown in [Figure 1](#).

**Figure 1. Active-Active Configuration Before Failover**



**Chassis A:**

On Chassis A, two failover groups are associated with the fw1 VAP group.

Failover group F1 has a configured VRRP priority of 250 compared to a configured VRRP priority of 249 for the associated F1 failover group on Chassis B. As long as these priorities remain unchanged, failover group F1 on Chassis A is designated as master and failover group F1 on Chassis B is backup.

Failover group F2 has a configured VRRP priority of 249 compared to a configured VRRP priority of 250 for the associated F2 failover group on Chassis B. As long as these priorities remain unchanged, failover group F2 on Chassis A is designated as backup to failover group F2 on Chassis B.

Chassis A processes traffic through the fw1 VAP group until a failure occurs that lowers the VRRP priority of the F1 failover group below the VRRP priority of the associated F1 failover group on Chassis B. At that time, a failover occurs and the traffic that was being processed by the fw1 VAP group begins to be processed by the fw2 VAP group.

**Chassis B:**

On Chassis B, two failover groups are associated with the fw2 VAP group.

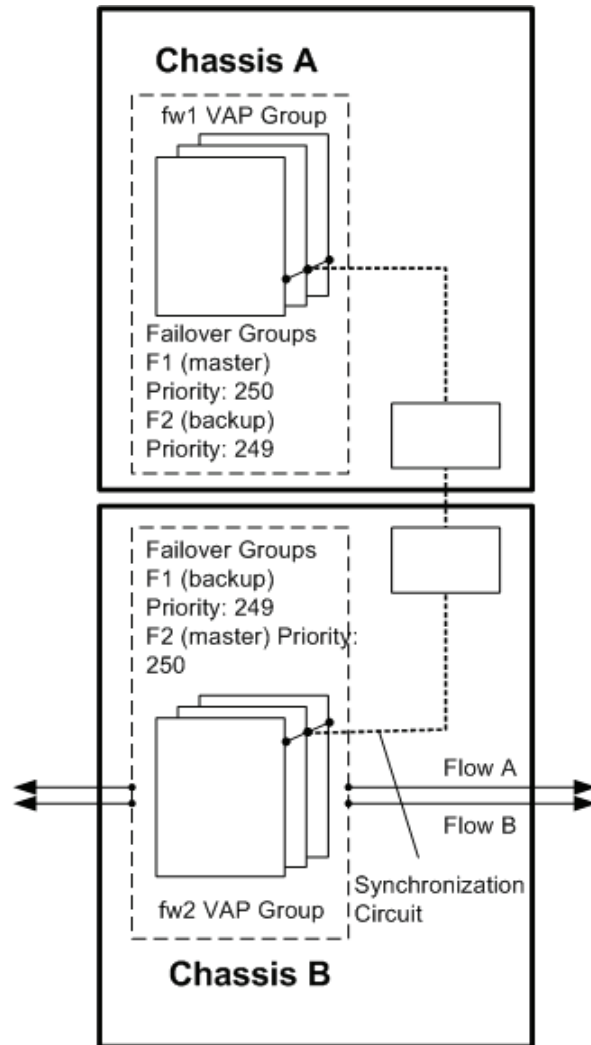
Failover group F1 has a configured VRRP priority of 249 compared to a configured VRRP priority of 250 for the associated F1 failover group on Chassis A. As long as these priorities remain unchanged, failover group F1 on Chassis A is designated as master and F1 on Chassis B is backup.

Failover group F2 has a configured VRRP priority of 250 compared to a configured VRRP priority of 249 for the associated F2 failover group on Chassis A. As long as these priorities remain unchanged, failover group F2 on Chassis B is designated as master and failover group F2 on Chassis A is backup.

Chassis B processes traffic through the fw2 VAP group until a failure occurs that lowers the VRRP priority of the F2 failover group below the VRRP priority of the associated F2 failover group on Chassis A. At that time, a failover occurs and the traffic that was being processed by the fw2 VAP group begins to be processed by the fw1 VAP group.

**NOTE:** When any failure occurs, the actual VRRP priority of the failover groups is compared, not the configured VRRP priority. If both chassis have experienced failures, the failover group with the higher actual priority is designated as master.

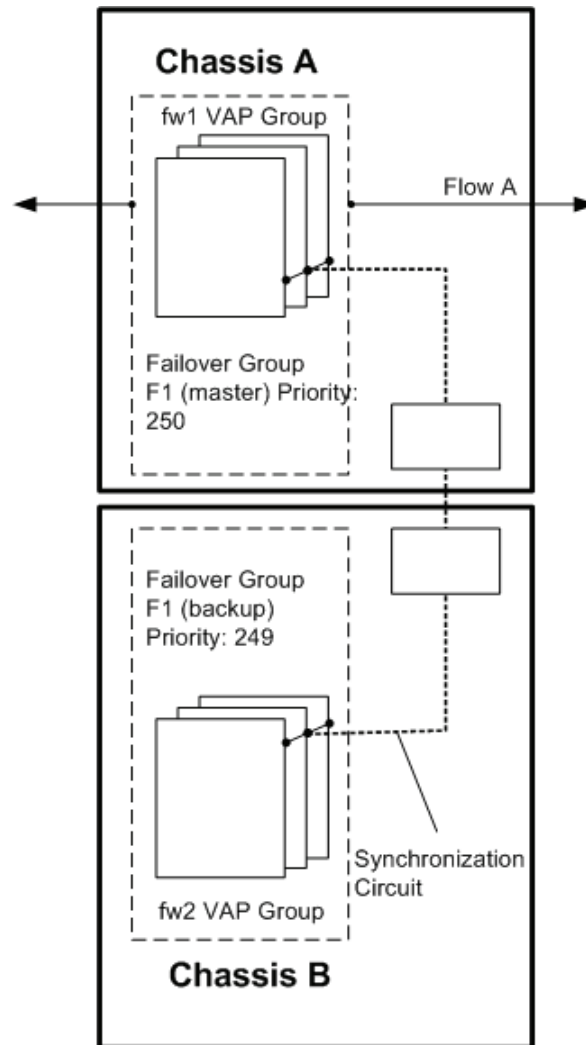
Figure 2. Active-Active Configuration After Failover to Chassis B



# Configuring an Active-Standby High Availability System

In an Active-Standby configuration, one system processes traffic and the other is in standby mode. Each has a failover group configured. The basic configuration, is shown in [Figure 1](#).

**Figure 3. Active-Standby Configuration Before Failover**



## Chassis A:

On Chassis A, one failover group is associated with the fw1 VAP group.

Failover group F1 has a configured VRRP priority of 250 compared to a configured VRRP priority of 249 for the associated F1 failover group on Chassis B. As long as these priorities remain unchanged, failover group F1 on Chassis A is designated as master and failover group F1 on Chassis B is backup.

Chassis A processes traffic through the fw1 VAP group until a failure occurs that lowers the VRRP priority of the F1 failover group below the VRRP priority of the associated F1 failover group on Chassis B. At that time, a failover occurs and the traffic that was being processed by the fw1 VAP group begins to be processed by the fw2 VAP group.

**Chassis B:**

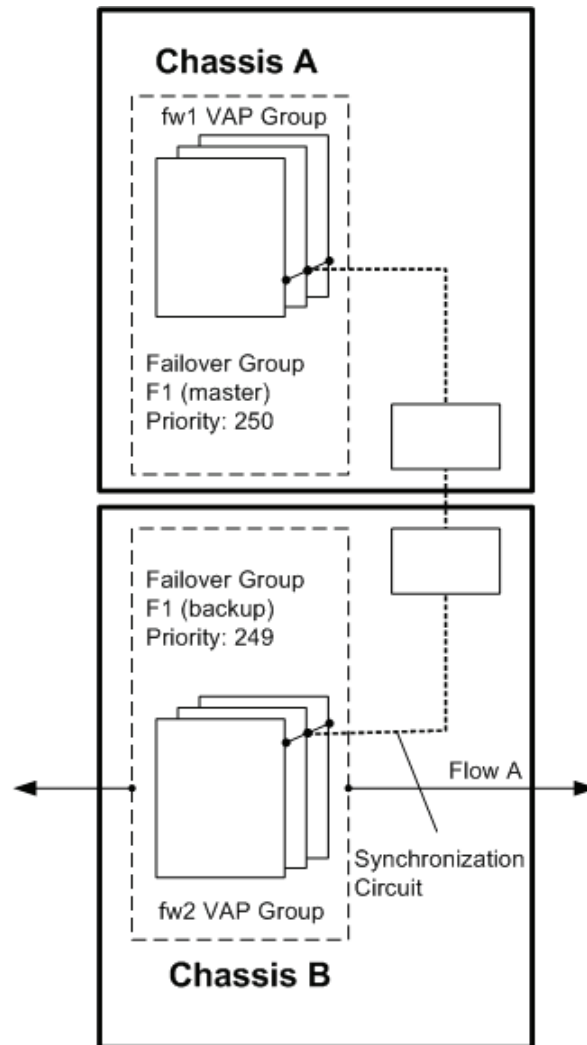
On Chassis B, one failover group is associated with the fw2 VAP group.

Failover group F1 has a configured VRRP priority of 249 compared to a configured VRRP priority of 250 for the associated F1 failover group on Chassis A. As long as these priorities remain unchanged, failover group F1 on Chassis A is designated as master and F1 on Chassis B is backup.

Chassis B does not process traffic until a failure occurs that lowers the VRRP priority of the F1 failover group on Chassis A below the VRRP priority of the associated F1 failover group on Chassis B. At that time, a failover occurs and the traffic that was being processed by the fw1 VAP group begins to be processed by the fw2 VAP group.

**NOTE:** When any failure occurs, the actual VRRP priority of the failover groups is compared, not the configured VRRP priority. If both chassis have experienced failures, the failover group with the higher actual priority is designated as master.

Figure 4. Active-Standby Configuration After Failover



---

# Active-Standby VRRP Dual-box, High Availability Configuration

This chapter provides detailed information about setting up two X-Series platforms in an Active-Standby configuration. The Active platform processes traffic while the Standby platform is idle, ready to take over if the Active chassis experiences a problem.

## Chassis Hardware Configurations

This chapter assumes the following:

**Chassis 1** has the following hardware configuration:

- Internal network: 1.1.45.0/16 (System ID 45)
- Two CPMs
  - ◆ CP1 internal IP address: 1.1.45.20 (Primary)
  - ◆ CP2 internal IP address: 1.1.45.21 (Secondary)
- Four NPMs (NP1, NP2, NP3, and NP4)
- Eight APMs (AP3, AP4, AP5, . . . and AP10)
- Management Interface IP addresses:
  - ◆ 192.168.50.45 (Mgmt 13/1)
  - ◆ 192.168.51.55 (Mgmt 13/2)
  - ◆ 192.168.50.65 (Mgmt 14/1)
  - ◆ 192.168.51.75 (Mgmt 14/2)

**NOTE:** By default, CPM management interfaces are not configured but can be configured if desired. The examples in this document include management interface information for illustration purposes.

**Chassis 2** has the following hardware configuration:

- Internal network: 1.1.46.0/16 (System ID 46)
- Two CPMs
  - ◆ CP1 internal IP address: 1.1.46.20 (Primary)
  - ◆ CP2 internal IP address: 1.1.46.21 (Secondary)
- Four NPMs (NP1, NP2, NP3, and NP4)
- Eight APMs (AP3, AP4, AP5, . . . and AP10)

- Management Interface IP addresses:
  - ◆ 192.168.50.46 (Mgmt 13/1)
  - ◆ 192.168.51.56 (Mgmt 13/2)
  - ◆ 192.168.50.66 (Mgmt 14/1)
  - ◆ 192.168.51.76 (Mgmt 14/2)

**NOTE:** By default, CPM management interfaces are not configured but can be configured if desired. The examples in this document include management interface information for illustration purposes.

## Assumptions

This document assumes that:

- You have set up your two chassis for basic operation.
- You have installed a Check Point firewall application.

For instructions on how to perform these tasks, see the list of documents in [Software Documentation](#) on page 5.

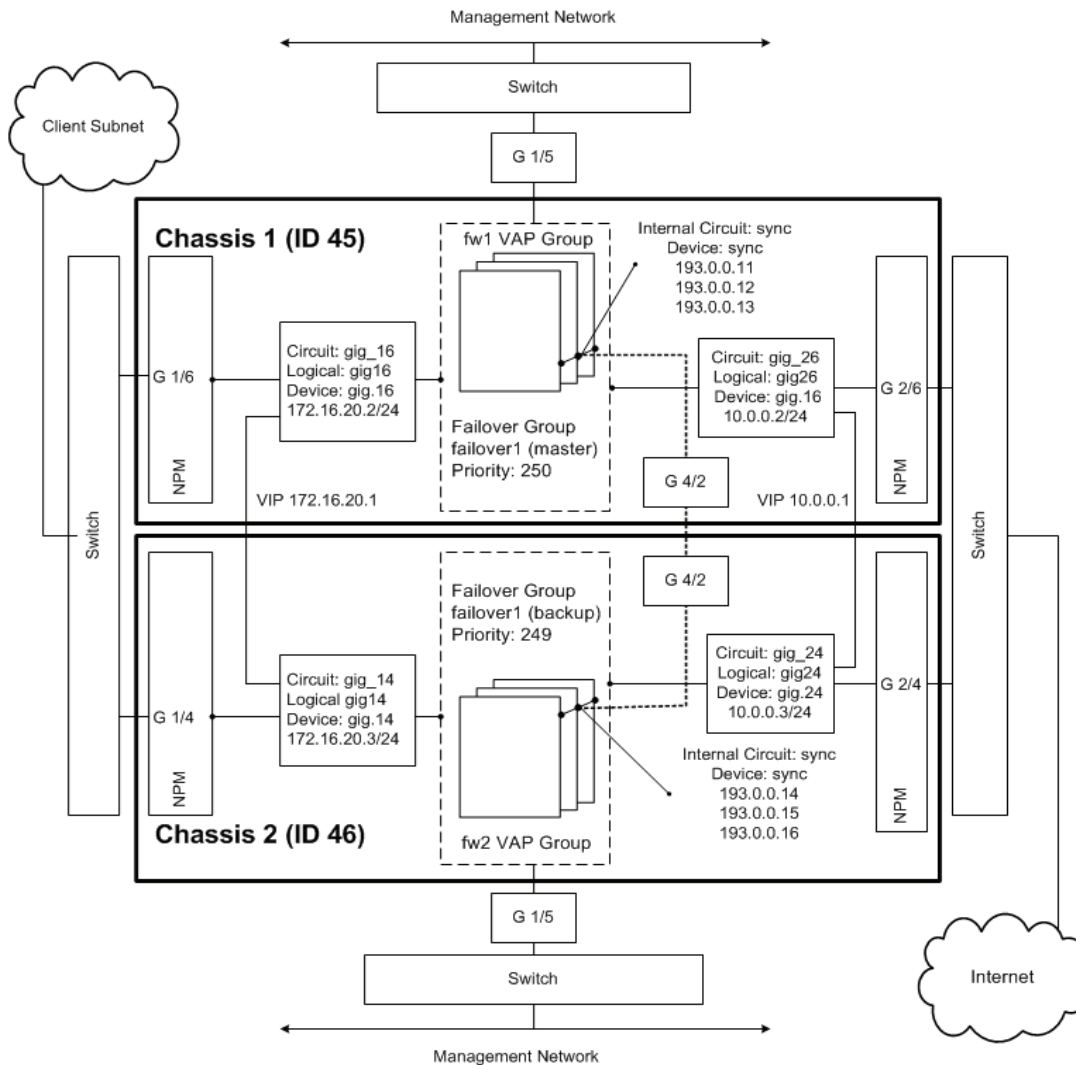
## Active-Standby Configuration

This section describes how to configure the two chassis for Active-Standby VRRP Dual-box High Availability operation. See:

- [System Diagram \(Active-Standby\)](#) on page 19
- [Configuring Chassis 1](#) on page 19
- [Configuring Chassis 2](#) on page 27

# System Diagram (Active-Standby)

This diagram illustrates the goal of the configuration steps in this chapter.



## 1.0 Configuring Chassis 1

On Chassis 1, perform these tasks:

- [Configuring System-wide Parameters \(Chassis 1\)](#)
- [Configuring Failover Group 1 \(Chassis 1\)](#)

### 1.1 Configuring System-wide Parameters (Chassis 1)

#### 1.1.1 Local System Identifier

Configure the local system-identifier on Chassis 1.

**NOTE:** When configuring multiple chassis for high availability, a unique system ID must be assigned to each chassis. If both chassis are configured with the same ID, you run the risk of having identical MAC addresses on any given circuit. This configuration is **not** supported or recommended.

If your X-Series Platforms do not have unique system IDs assigned, use the following command to define a system ID. The valid range is from 1-255.

```
CBS# configure system-identifier 45
```

**NOTE:** After you configure a system-identifier, you must use the `reload all` command to activate the identifier.

### 1.1.2 Remote System Identifier

**NOTE:** If you connect Chassis 1 to Chassis 2 using only the High Availability (HA) port, you do not need to configure the remote system identifier. The two chassis use the HA port to exchange system-id information. However, if you do this, the HA port connection represents a single point of failure. Crossbeam recommends that you connect the HA ports and at least one pair of management interfaces.

Configure the remote system ID and IP address using the `configure remote-box` command.

**NOTE:** The `remote-box` command requires that you have interconnected CPMs on the two chassis. The IP address that you specify depends on how you have interconnected the CPMs. If you use the High Availability port, specify the **internal** IP address (1.1.46.20) associated with the remote Primary CPM (obtained by running `show-internal-ip` on the remote chassis). If you use either or both of the management ports, specify the **external** IP address(es) associated with the port(s). You can specify all three addresses within a single `configure remote-box` command. Crossbeam recommends that you connect the CPMs using separate network broadcast domains for each pair of ports (the HA ports, the Management 1 ports, and the Management 2 ports). Connecting the ports directly can lead to scenarios that do not provide full redundancy.

```
CBS# configure remote-box 46 1.1.46.20 192.168.50.46 192.168.51.56
CBS(conf-remote-box) # end
```

**NOTE:** The example `configure remote-box` command specifies only the IP addresses of the management 1 and 2 interfaces for one of the CPMs on chassis 2. If you have connected the management interfaces of the other CPM on Chassis 2, you can add the IP addresses for those interfaces.

### 1.1.3 Configure the Synchronization Circuit

Configure a synchronization circuit between VAP Group fw1 on Chassis 1 and fw2 on Chassis 2 so that the two VAP Groups act as one Check Point Cluster with 6 members.

**NOTE:** This step must be performed **after** you configure the system-id, because the system-id affects the MAC selection and configuration of every circuit that gets created.

**NOTE:** When you configure a circuit, you can either assign a circuit-id number (1-4095) or let XOS assign the next available circuit-id number. This example illustrates the manual configuration method.

Enter these commands:

```
CBS# configure circuit sync circuit-id 100
CBS(conf-cct) # device-name sync
CBS(conf-cct) # internal
CBS(conf-cct) # vap-group fw1
CBS(conf-cct-vapgroup) # ip 193.0.0.11/24 increment-per-vap 193.0.0.13
CBS(conf-cct-vapgroup) # end
CBS#
```

```

CBS# configure interface gigabitethernet 4/2
CBS(conf-interface-gig)# logical sync
CBS(confi-gig-logical)# circuit sync
CBS(intf-logical)# end
CBS#

```

### 1.1.4 Configure the Traffic Circuits

Configure the two circuits that convey traffic to and from the fw1 VAP group.

**NOTE:** When you configure a circuit, you can either assign a circuit-id number (1-4095) or let XOS assign the next available circuit-id number. This example illustrates the manual configuration method.

```

CBS# configure circuit gig_16 circuit-id 106 device-name gig.16
CBS# configure circuit gig_16 vap-group fw1
CBS(conf-cct-vapgroup)# ip 172.16.20.2/24
CBS(conf-cct-vapgroup-ip)# enable
CBS(conf-cct-vapgroup-ip)# end
CBS# configure circuit gig_26 circuit-id 206 device-name gig.26
CBS# configure circuit gig_26 vap-group fw1
CBS(conf-cct-vapgroup)# ip 10.0.0.2/24
CBS(conf-cct-vapgroup-ip)# enable
CBS(conf-cct-vapgroup-ip)# end
CBS#

```

### 1.1.5 Configure the Traffic Interfaces

Configure the interfaces through which traffic flows to and from the fw1 VAP group.

```

CBS# configure interface gigabitethernet 1/6
CBS(conf-intf-gig)# logical gig16
CBS(intf-gig-logical)# circuit gig_16
CBS(intr-gig-log-cct)# end
CBS#
CBS# configure interface gigabitethernet 2/6
CBS(conf-intf-gig)# logical gig26
CBS(intf-gig-logical)# circuit gig_26
CBS(intr-gig-log-cct)# end

```

## 1.2 Configuring Failover Group 1 (Chassis 1)

### 1.2.1 VRRP Failover Group

Create the failover group by assigning it a name (`failover1`) and a failover group ID. The failover group ID is different than the system identifier, configured earlier. The ID must be unique on this chassis, and must be the same on its counterpart failover group on the remote chassis (Chassis 2).

The `failover1` group acts as the master group on Chassis 1. The counterpart failover group on Chassis 2 is also called `failover1` and has the same group ID (1). The two groups have different `priority` values.

```

CBS# configure vrrp failover-group failover1 failover-group-id 1
CBS(conf-vrrp-group)#

```

## 1.2.2 VRRP Priority

For proper operation, the VRRP `priority` value of the two associated failover-groups must be different on Chassis 1 and Chassis 2; during normal operations, the failover group with the higher priority is the master. Certain events such as an interface failure or a change in VAP group member count can be configured to decrement the VRRP priority of the failover-group. Failover occurs when the VRRP `priority` value of one failover group drops below the priority of the failover group on the other chassis. VRRP `priority` values range from 1 to 255, and the default is 100.

```
CBS(conf-vrrp-group) # priority 250
CBS(conf-vrrp-group) # exit
```

## 1.2.3 Virtual Router on each Traffic Circuit

Create a virtual router on each traffic circuit that is attached to the **fw1** VAP group. A virtual router is assigned a virtual IP addresses that is used to configure VRRP and, optionally, next hop health check. This section describes the configuration of two virtual routers, one for each of the two circuits (`gig_16` and `gig_26`) that are associated with the **fw1** VAP group.

### Configuring the virtual-router for the `gig_16` circuit

1. To create the virtual router for the first circuit, enter this command:

```
CBS(conf-vrrp-group) # virtual-router vrrp-id 10 circuit gig_16
CBS(conf-vrrp-failover-vr) #
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr) # priority-delta 2
CBS(conf-vrrp-failover-vr) #
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 1 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 2, so the failover group on Chassis 2 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router

```
CBS(conf-vrrp-failover-vr) # mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr) #
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router

```
CBS(conf-vrrp-failover-vr) # vap-group fw1
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** Before you map the virtual router to the VAP group, the circuit must have been mapped to the VAP group.

Mapping the VAP group to the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 172.16.20.1/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

## Configuring the virtual-router for the gig\_26 circuit

1. To create the virtual router for the second circuit, enter this command:

```
CBS(conf-vrrp-group)# virtual-router vrrp-id 11 circuit gig_26  
CBS(conf-vrrp-failover-vr)#
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr)# priority-delta 2  
CBS(conf-vrrp-failover-vr)#
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 1 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 2, so the failover group on Chassis 2 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router

```
CBS(conf-vrrp-failover-vr)# mac-usage vrrp-mac  
CBS(conf-vrrp-failover-vr)#
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router

```
CBS(conf-vrrp-failover-vr)# vap-group fw1  
CBS(conf-vrrp-vr-vapgroup)#
```

**NOTE:** The circuit must already be mapped to the VAP group.

Specifying the VAP group of the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup)# virtual-ip 10.0.0.1/24  
CBS(conf-vrrp-vr-vapgroup)# end  
CBS#
```

### 1.2.4 Enable VRRP on the VAP Group

VRRP monitors the `fw1` VAP group for failure of individual VAPs. By setting the `active-vap-threshold` and the `priority-delta`, individual VAP failures can decrement VRRP priority and cause a comparison with the VRRP priority on the remote chassis. Failover occurs when the `priority-delta` value decrements the `priority` value of the master failover group below the `priority` value of the associated failover group on the remote chassis. In our configuration, the `priority` value for the master failover group is 250 and the `priority` value for the backup group is 249. A `priority-delta` of 2 is used for each of the VAP groups.

To configure the `fw1` VAP group for failover:

1. Enable VRRP on the `fw1` VAP group.

```
CBS# configure vrrp vap-group fw1  
CBS(conf-vrrp-vap-group)#
```

2. Assign the VRRP enabled `fw1` VAP group to a failover group list (required).

```
CBS(conf-vrrp-vap-group)# failover-group-list failover1  
CBS(conf-vrrp-vap-group)#
```

3. Optionally, set the hold down timer.

Configure the `hold-down-timer` to 120, which forces the VAP group to wait for two minutes while the application fully boots. This wait prevents the failover group from becoming VRRP master before the application is fully active.

```
CBS(conf-vrrp-vap-group) # hold-down-timer 120
CBS(conf-vrrp-vap-group) #
```

4. Optionally, set the active VAP threshold and return to the main CLI context.

The `active-vap-threshold` monitors the number of active VAPs in the VAP group. If the number of active VAPs drops below the threshold, the failover group's `priority` value is decremented by the `priority-delta` (defined in [Step 5](#)) and a comparison is done between the priorities of the failover groups on the two chassis. Failover occurs when the `priority-delta` decrements the `priority` value of the master failover group below the `priority` value of the backup group.

```
CBS(conf-vrrp vap-group) # active-vap-threshold 3
CBS(conf-vrrp vap-group) # end
CBS#
```

5. Set the `priority-delta` value for the VAP group (optional).

Assign a `priority-delta` to the VAP group. VRRP decrements the priority of the failover group whenever the number of active VAPs falls below the `active-vap-threshold`.

The `priority-delta` value can be any number between 1 and 255 and the default value is 1. When the VAP returns to the Active state, the `priority-delta` value is added back to the `priority` value.

```
CBS(conf-vrrp-vap-group) # priority-delta 2
CBS(conf-vrrp-vap-group) #
```

## 1.3 Optionally Configuring OSPF

If your network is configured to use the Open Shortest Path First (OSPF) protocol, you can incorporate OSPF into your VRRP configuration.

**NOTE:** To configure OSPF, you must first install the Crossbeam Routing Software (RSW).

When a failover occurs from one failover group to another, you want traffic to be rerouted from the failed group to the one that is now active. To ensure that this happens, you can increase the `ospf-cost-increment` value associated with the circuit on the first failover group. The new value is propagated to all local routers, increasing the OSPF cost of the circuit so that it is no longer part of the preferred route. When the original failover group resumes master status, the OSPF cost is readjusted to the originally configured value.

**NOTE:** Configure the `ospf-cost-increment` only on the master failover group.

To include OSPF cost in the configuration, perform these steps:

1. Configure these parameters on the master failover group (failover1):

```
CBS# configure vrrp failover-group failover1
CBS(conf-vrrp-group) # ospf-cost-increment circuit gig_16 5
CBS(conf-vrrp-group) # ospf-cost-increment circuit gig_26 5
```

2. On the VAP group that is associated with the master failover group, start the ospf routing protocol.

```
CBS# configure routing-protocol ospf vap-group fw1 start
```

## 1.4 Preemption

Typically, after a failover has occurred from one failover group to another, you want the failover group that is now master to remain in that role, even after the problem that caused the failover has been resolved.

By default, preemption is turned off, and Crossbeam recommends that you do not configure preemption for failover groups.

However, if you want the original failover group to resume the role of master, you can turn preemption on using this command:

```
CBS(conf-vrrp-group) # preemption
```

## 1.5 Management Circuit

If you intend to run Check Point software on the X-Series chassis and you intend to use a Check Point management station, do not configure the X-Series chassis management circuits as part of any failover group. If you do, the Check Point management station cannot access the management circuit.

## Verifying Your Configuration

This section contains the output of several commands that show the configured status of Chassis 1. To check your progress throughout the configuration process, open a second CLI window, log in to the CPM and enter one of the commands.

### Output of show running-config on Chassis 1

```
CBS# show running-config
#
vap-group fw1 xslinux_v5
  vap-count 3
  max-load-count 3
  ap-list ap1 ap2 ap3 ap4 ap5 ap6 ap7 ap8 ap9 ap10
  load-balance-vap-list 1 2 3 4 5 6 7 8 9 10
  ip-flow-rule loadbalance1
    action load-balance
    activate
#
circuit gig_16 circuit-id 106
  device-name gig.16
  vap-group fw1
  ip 172.16.20.2/24 172.16.20.255
circuit gig_26 circuit-id 206
  device-name gig.26
  vap-group fw1
  ip 10.0.0.2/24 10.0.0.255
#
interface gigabitethernet 1/6
  logical gig16
  circuit gig_16
interface gigabitethernet 2/6
  logical gig26
  circuit gig_26
#
vrrp failover-group failover1 failover-group-id 1
  priority 250
  virtual-router vrrp-id 10 circuit gig_16
    priority-delta 2
    vap-group fw1
    virtual-ip 172.16.20.1/24 172.16.20.255
  virtual-router vrrp-id 11 circuit gig_26
    priority-delta 2
    vap-group fw1
```

```

        virtual-ip 10.0.0.1/24 10.0.0.255
#
vrrp vap-group fw1
    failover-group-list failover1
    hold-down-timer 120
    priority-delta 2
#
management gigabitethernet 13/1
    ip-addr 192.168.50.45/24 192.168.50.255
    enable
    access-list 1 input
    access-list 2 output
management gigabitethernet 13/2
    ip-addr 192.168.51.55/24 192.168.51.255
    enable
    access-list 1 input
    access-list 2 output
management gigabitethernet 14/1
    ip-addr 192.168.50.65/24 192.168.50.255
    enable
    access-list 1 input
    access-list 2 output
management gigabitethernet 14/2
    ip-addr 192.168.51.75/24 192.168.51.255
    enable
    access-list 1 input
    access-list 2 output

```

### Output of show vrrp on Chassis 1

```

CBS# show vrrp
Priority is Actual/Configured
FG-ID  Priority  Status  Preempt  Master Sys ID  Master Priority
1      250/250    Master  off      45           250
(1 row)

```

### Output of show vrrp detail-status on Chassis 1

```

CBS# show vrrp detail-status
FG_ID  Status  Priority  Delta  Type  Component
1      Master  250/250  2      vr   gig_16/10/5
1      Master  250/250  2      vr   gig_26/10/5
1      Master  250/250  2      vg   fw1

```

**NOTE:** The Component column shows 5 as the last digit only if you configured the OSPF cost increment as described in [Optionally Configuring OSPF](#) on page 24. If you did not configure the OSPF cost increment, these values would be 0 (zero).

### Output of show remote-box on Chassis 1

```

CBS# show remote-box
System ID      : 46
First IP Address : 1.1.46.20
Second IP Address : 192.168.50.46
Third IP Address  : 192.168.51.56
Fourth IP Address : 0.0.0.0
Fifth IP Address  : 0.0.0.0
Active IP Address : 1.1.46.20
(1 row)

```

## 2.0 Configuring Chassis 2

On Chassis 2, perform these tasks:

- [Configuring System-wide Parameters \(Chassis 2\)](#)
- [Configuring Failover Group 1 \(Chassis 2\)](#)

### 2.1 Configuring System-wide Parameters (Chassis 2)

#### 2.1.1 Local System Identifier

Configure the local system-identifier on Chassis 2.

**IMPORTANT:** When configuring multiple chassis for high availability, a unique system ID must be assigned to each chassis. If both chassis are configured with the same ID, you run the risk of having identical MAC addresses on any given circuit. This configuration is **not** supported or recommended.

If your X-Series Platforms do not have system IDs assigned, use the following command to define a system ID. The valid range is from 1-255.

```
CBS# configure system-identifier 46
```

**NOTE:** After you configure a system-identifier, you must use the `reload all` command to activate the identifier.

#### 2.1.2 Remote System Identifier

**NOTE:** If you connect Chassis 1 to Chassis 2 using only the High Availability (HA) port, you do not need to configure the remote system identifier. The two chassis use the HA port to exchange system-id information. However, if you do this, the HA port connection represents a single point of failure. Crossbeam recommends that you connect the HA ports and at least one pair of management interfaces.

Configure the remote system ID and IP address using the `configure remote-box` command.

**NOTE:** The `remote-box` command requires that you have interconnected CPMs on the two chassis. The IP address that you specify depends on how you have interconnected the CPMs. If you use the High Availability port, specify the **internal** IP address (1.1.45.20) associated with the remote Primary CPM (obtained by running `show-internal-ip` on the remote chassis). If you use either or both of the management ports, specify the **external** IP address(es) associated with the port(s). You can specify all three addresses within a single `configure remote-box` command. Crossbeam recommends that you connect the CPMs using separate network broadcast domains for each pair of ports (the HA ports, the Management 1 ports, and the Management 2 ports). Connecting the ports directly can lead to scenarios that do not provide full redundancy.

```
CBS# configure remote-box 45 1.1.45.20 192.168.50.45 192.168.51.55
CBS(conf-remote-box) # end
```

**NOTE:** The example `configure remote-box` command specifies only the IP addresses of the management 1 and 2 interfaces for one of the CPMs on chassis 1. If you have connected the management interfaces of the other CPM on Chassis 1, you can add the IP addresses for those interfaces.

#### 2.1.3 Configuring the Synchronization Circuit

Configure a synchronization circuit between VAP Group **fw1** on Chassis 1 and **fw2** on Chassis 2 so that the two VAP Groups act as one Check Point Cluster with 6 members.

**NOTE:** This step must be performed **after** you configure the system-id, because the system-id affects the MAC selection and configuration of every circuit that gets created.

Enter these commands:

```
CBS# configure circuit sync circuit-id 100
CBS(conf-cct) # device-name sync
CBS(conf-cct) # internal
CBS(conf-cct) # vap-group fw2
CBS(conf-cct-vapgroup) # ip 193.0.0.14/24 193.0.0.255 increment-per-vap
193.0.0.16
CBS(conf-cct-vapgroup) # end
CBS#
CBS# configure interface gigabitethernet 4/2
CBS(conf-interface-gig) # logical sync
CBS(confi-gig-logical) # circuit sync
CBS(intf-logical) # end
CBS#
```

## 2.1.4 Configure the Traffic Circuits

Configure the two circuits that convey traffic to and from the fw2 VAP group.

**NOTE:** When you configure a circuit, you can either assign a circuit-id number (1-4095) or let XOS assign the next available circuit-id number. This example illustrates the manual configuration method.

```
CBS# configure circuit gig_14 circuit-id 104 device-name gig.14
CBS# configure circuit gig_14 vap-group fw2
CBS(conf-cct-vapgroup) # ip 172.16.20.3/24
CBS(conf-cct-vapgroup-ip) # enable
CBS(conf-cct-vapgroup-ip) # end
CBS# configure circuit gig_24 circuit-id 204 device-name gig.24
CBS# configure circuit gig_24 vap-group fw2
CBS(conf-cct-vapgroup) # ip 10.0.0.3/24
CBS(conf-cct-vapgroup-ip) # enable
CBS(conf-cct-vapgroup-ip) # end
CBS#
```

## 2.1.5 Configure the Traffic Interfaces

Configure the interfaces through which traffic flows to and from the fw1 VAP group.

```
CBS# configure interface gigabitethernet 1/4
CBS(conf-intf-gig) # logical gig14
CBS(intf-gig-logical) # circuit gig_14
CBS(intr-gig-log-cct) # end
CBS#
CBS# configure interface gigabitethernet 2/4
CBS(conf-intf-gig) # logical gig24
CBS(intf-gig-logical) # circuit gig_24
CBS(intr-gig-log-cct) # end
```

## 2.2 Configuring Failover Group 1 (Chassis 2)

### 2.2.1 VRRP Failover Group

Create the failover group by assigning it a name (`failover1`) and a failover group ID. The failover group ID is different than the system identifier, configured earlier. The ID must be unique on this chassis, and must be the same on its counterpart failover group on the remote chassis (Chassis 1).

The failover1 group on Chassis 2 acts as the backup group to the failover1 group on Chassis 1. The two groups have different `priority` values.

```
CBS# configure vrrp failover-group failover1 failover-group-id 1
CBS(conf-vrrp-group) #
```

## 2.2.2 VRRP Priority

For proper operation, the VRRP `priority` value of the two associated failover-groups must be different on Chassis 1 and Chassis 2; during normal operations, the failover group with the higher priority is the master. Certain events such as an interface failure or a change in VAP group member count can be configured to decrement the VRRP priority of the failover-group. Failover occurs when the VRRP `priority` value of one failover group drops below the priority of the failover group on the other chassis. VRRP `priority` values range from 1 to 255, and the default is 100.

```
CBS(conf-vrrp-group) # priority 249
CBS(conf-vrrp-group) #
```

## 2.2.3 Virtual Router on each Traffic Circuit

Create virtual routers on each traffic circuit that is attached to the **fw2** VAP group. A virtual router is assigned a virtual IP addresses that is used to configure VRRP and, optionally, next hop health check. This section describes the configuration of two virtual routers, one for each of the two circuits (`gig_14` and `gig_24`) that are associated with the **fw2** VAP group.

## Configuring the virtual-router for the `gig_14` circuit

1. To create the virtual router for the first circuit, enter this command:

```
CBS(conf-vrrp-group) # virtual-router vrrp-id 10 circuit gig_14
CBS(conf-vrrp-failover-vr) #
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr) # priority-delta 2
CBS(conf-vrrp-failover-vr) #
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 2 becomes 247, which remains less than the `priority` value of the associated failover group on Chassis 1, so the failover group on Chassis 1 remains master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router

```
CBS(conf-vrrp-failover-vr) # mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr) #
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router

```
CBS(conf-vrrp-failover-vr) # vap-group fw2
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** Before you map the virtual router to the VAP group, the circuit must have been mapped to the VAP group.

Mapping the VAP group to the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 172.26.20.1/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

## Configuring the virtual-router for the gig\_24 circuit

1. To create the virtual router for the second circuit, enter this command:

```
CBS(conf-vrrp-group) # virtual-router vrrp-id 11 circuit gig_24
CBS(conf-vrrp-failover-vr) #
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr) # priority-delta 2
CBS(conf-vrrp-failover-vr) #
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 2 becomes 247, which remains less than the `priority` value of the associated failover group on Chassis 1, so the failover group on Chassis 1 remains master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router

```
CBS(conf-vrrp-failover-vr) # mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr) #
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router

```
CBS(conf-vrrp-failover-vr) # vap-group fw2
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** The circuit must already be mapped to the VAP group.

Specifying the VAP group of the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 10.0.0.1/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

## 2.2.4 Enable VRRP on the VAP Group

VRRP monitors the **fw2** VAP group for failure of individual VAPs. By setting the `active-vap-threshold` and the `priority-delta`, individual VAP failures can decrement VRRP priority and cause a comparison with the VRRP priority on the remote chassis. Failover occurs when the `priority-delta` decrements the `priority` value of the master failover group below the `priority` value of the associated failover group on the remote chassis. In the example configuration, priority for the master failover group is 250 and priority for the backup group is 249. A `priority-delta` of 2 is used for each of the VAP groups.

To configure the **fw2** VAP group for failover:

1. Enable VRRP on the **fw2** VAP group.

```
CBS# configure vrrp vap-group fw2  
CBS(conf-vrrp-vap-group) #
```

2. Assign the VRRP enabled **fw2** VAP group to a failover group list (required).

```
CBS(conf-vrrp-vap-group) # failover-group-list failover1  
CBS(conf-vrrp-vap-group) #
```

3. Optionally, set the hold down timer.

Configure the `hold-down-timer` to 120, which forces the VAP group to wait for two minutes while the application fully boots. This wait prevents the failover group from becoming VRRP master before the application is fully active.

```
CBS(conf-vrrp-vap-group) # hold-down-timer 120  
CBS(conf-vrrp-vap-group) #
```

4. Optionally, set the active VAP threshold and return to the main CLI context.

The `active-vap-threshold` monitors the number of active VAPs in the VAP group. If the number of active VAPs drops below the threshold, the failover group's `priority` value is decremented by the `priority-delta` (defined in [Step 5](#)) and a comparison is done between the priorities of the failover groups on the two chassis. Failover occurs when the `priority-delta` decrements the `priority` value of the master failover group below the `priority` value of the backup group.

```
CBS(conf-vrrp vap-group) # active-vap-threshold 3  
CBS(conf-vrrp vap-group) # end  
CBS#
```

5. Set the `priority-delta` value for the VAP group (optional).

Assign a `priority-delta` to the VAP group. VRRP decrements the priority of the failover group whenever the number of active VAPs falls below the `active-vap-threshold`.

The `priority-delta` value can be any number between 1 and 255 and the default value is 1. When the VAP returns to the Active state, the `priority-delta` value is added back to the `priority` value.

```
CBS(conf-vrrp-vap-group) # priority-delta 2  
CBS(conf-vrrp-vap-group) #
```

## 2.3 Preemption

Typically, after a failover has occurred from one failover group to another, you want the failover group that is now master to remain in that role, even after the problem that caused the failover has been resolved.

By default, preemption is turned off, and Crossbeam recommends that you do not configure preemption for failover groups.

However, if you want the original failover group to resume the role of master, you can turn preemption on using this command:

```
CBS(conf-vrrp-group) # preemption
```

## 2.4 Management Circuit

If you intend to run Check Point software on the X-Series chassis and you intend to use a Check Point management station, do not configure the X-Series chassis management circuits as part of any failover group. If you do, the Check Point management station cannot access the management circuit.

### Verifying Your Configuration

This section contains the output of several commands that show the configured status of Chassis 1. To check your progress throughout the configuration process, open a second CLI window, log in to the CPM and enter one of the show commands.

#### Output of show running-config on Chassis 2

```
CBS# show running-config
#
vap-group fw2 xslinux_v5
  vap-count 3
  max-load-count 3
  ap-list ap1 ap2 ap3 ap4 ap5 ap6 ap7 ap8 ap9 ap10
  load-balance-vap-list 1 2 3 4 5 6 7 8 9 10
  ip-flow-rule loadbalance1
    action load-balance
    activate
#
circuit gig_14 circuit-id 104
  device-name gig.14
  vap-group fw2
    ip 172.16.20.3/24 172.16.20.255
circuit gig_24 circuit-id 204
  device-name gig.24
  vap-group fw2
    ip 10.0.0.3/24 10.0.0.255
#
interface gigabitethernet 1/4
  logical gig14
  circuit gig_14
interface gigabitethernet 2/4
  logical gig24
  circuit gig_24
#
vrrp failover-group failover1 failover-group-id 1
  priority 249
  virtual-router vrrp-id 10 circuit gig_14
    priority-delta 2
    mac-usage vrrp-mac
    vap-group fw2
      virtual-ip 172.16.20.1/24 172.16.20.255
  virtual-router vrrp-id 11 circuit gig_24
    priority-delta 2
    mac-usage vrrp-mac
    vap-group fw2
      virtual-ip 10.0.0.1/24 10.0.0.255
#
#
#
#
```

```

vrrp vap-group fw2
  failover-group-list failover1
  hold-down-timer 120
  priority-delta 2
#
management gigabitethernet 13/1
  ip-addr 192.168.50.46/24 192.168.50.255
  enable
  access-list 1 input
  access-list 2 output
management gigabitethernet 13/2
  ip-addr 192.168.51.56/24 192.168.51.255
  enable
  access-list 1 input
  access-list 2 output
management gigabitethernet 14/1
  ip-addr 192.168.50.66/24 192.168.50.255
  enable
  access-list 1 input
  access-list 2 output
management gigabitethernet 14/2
  ip-addr 192.168.51.76/24 192.168.51.255
  enable
  access-list 1 input
  access-list 2 output

```

### Output of show vrrp on Chassis 2

```

CBS# show vrrp
Priority is Actual/Configured
FG-ID  Priority  Status  Preempt  Master Sys ID  Master Priority
1      249/249    Backup off      45                250
(1 row)

```

### Output of show vrrp detail-status on Chassis 2

```

CBS# show vrrp detail-status
FG_ID  Status  Priority  Delta  Type  Component
1      Backup  249/249  2      vr    gig_14/10/0
1      Backup  249/249  2      vr    gig_24/11/0
1      Backup  249/249  2      vg    fw2

```

### Output of show remote-box on Chassis 2

```

CBS# show remote-box
System ID      : 45
First IP Address : 1.1.45.20
Second IP Address : 192.168.50.45
Third IP Address  : 192.168.51.55
Fourth IP Address : 0.0.0.0
Fifth IP Address  : 0.0.0.0
Active IP Address : 1.1.45.20
(1 row)

```



---

# Active-Active VRRP Dual-box, High Availability Configuration

This chapter provides detailed information about setting up two X-Series chassis in an Active-Active VRRP Dual-box High Availability configuration. Both chassis process traffic and, if either one experiences a problem, the other chassis assumes the workload of both.

## Chassis Hardware Configurations

This chapter assumes the following:

**Chassis 1** has the following hardware configuration:

- Internal network: 1.1.45.0/16 (System ID 45)
- Two CPMs
  - ◆ CP1 internal IP address: 1.1.45.20 (Primary)
  - ◆ CP2 internal IP address: 1.1.45.21 (Secondary)
- Four NPMs (NP1, NP2, NP3, and NP4)
- Eight APMs (AP3, AP4, AP5, . . . and AP10)
- CPM Management Interface IP addresses:
  - ◆ 192.168.50.45 (Mgmt 13/1)
  - ◆ 192.168.51.55 (Mgmt 13/2)
  - ◆ 192.168.50.65 (Mgmt 14/1)
  - ◆ 192.168.51.75 (Mgmt 14/2)

**NOTE:** By default, CPM management interfaces are not configured but can be configured if desired. The examples in this document include management interface information for illustration purposes.

**Chassis 2** has the following hardware configuration:

- Internal network: 1.1.46.0/16 (System ID 46)
- Two CPMs
  - ◆ CP1 internal IP address: 1.1.46.20 (Primary)
  - ◆ CP2 internal IP address: 1.1.46.21 (Secondary)
- Four NPMs (NP1, NP2, NP3, and NP4)
- Eight APMs (AP3, AP4, AP5, . . . and AP10)

- Management Interface IP addresses:
  - ◆ 192.168.50.46 (Mgmt 13/1)
  - ◆ 192.168.51.56 (Mgmt 13/2)
  - ◆ 192.168.50.66 (Mgmt 14/1)
  - ◆ 192.168.51.76 (Mgmt 14/2)

**NOTE:** By default, CPM management interfaces are not configured but can be configured if desired. The examples in this document include management interface information for illustration purposes.

## Assumptions

This document assumes that:

- You have set up your two chassis for basic operation.
- You have installed a Check Point firewall application.

For instructions on how to perform these tasks, see the list of documents in [Software Documentation](#) on page 5.

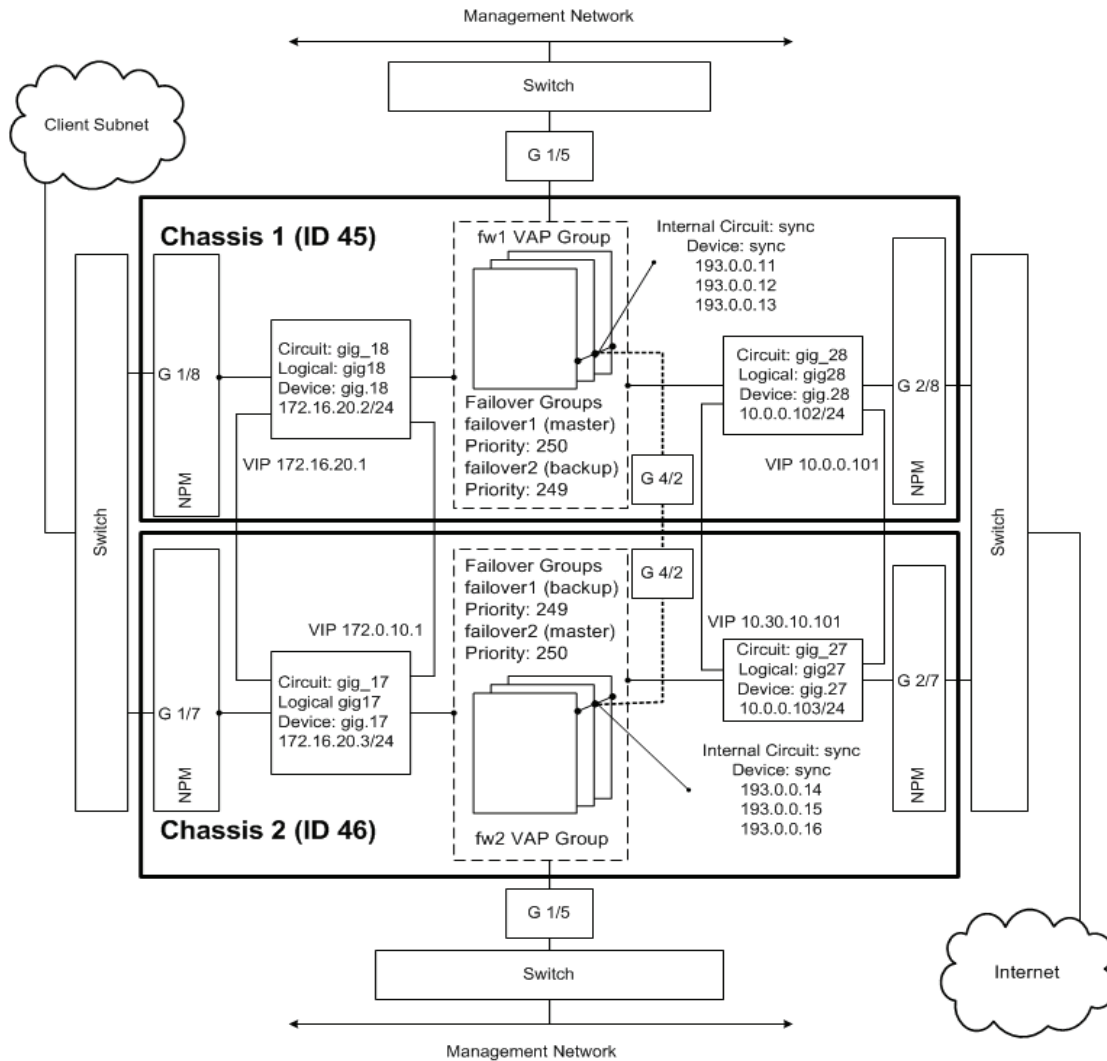
## Active-Active Configuration

This section describes how to configure the two chassis for Active-Active VRRP Dual-box High Availability operation. See:

- [System Diagram \(Active-Active\)](#) on page 37
- [Configuring Chassis 1](#) on page 38
- [Configuring Chassis 2](#) on page 50

# System Diagram (Active-Active)

This diagram illustrates the goal of the configuration steps in this chapter.



# 1.0 Configuring Chassis 1

On Chassis 1, perform these tasks:

- [Configuring System-wide Parameters \(Chassis 1\)](#)
- [Configuring Failover Group 1 \(Chassis 1\)](#)
- [Configuring Failover Group 2 \(Chassis 1\)](#)

## 1.1 Configuring System-wide Parameters (Chassis 1)

### 1.1.1 Local System Identifier

Configure the local system-identifier on Chassis 1.

**NOTE:** When configuring multiple chassis for high availability, a unique system ID must be assigned to each chassis. If both chassis are configured with the same ID, you run the risk of having identical MAC addresses on any given circuit. This configuration is **not** supported or recommended.

If your X-Series Platforms do not have unique system IDs assigned, use the following command to define a system ID. The valid range is from 1-255.

Command:

```
CBS# configure system-identifier 45
```

**NOTE:** After you configure a system-identifier, you must use the `reload all` command to activate the identifier.

### 1.1.2 Remote System Identifier

**NOTE:** If you connect Chassis 1 to Chassis 2 using only the High Availability (HA) port, you do not need to configure the remote system identifier. The two chassis use the HA port to exchange system-id information. However, if you do this, the HA port connection represents a single point of failure. Crossbeam recommends that you connect the HA ports and at least one pair of management interfaces.

Configure the remote system ID and IP address using the `configure remote-box` command.

**NOTE:** The `remote-box` command requires that you have interconnected CPMs on the two chassis. The IP address that you specify depends on how you have interconnected the CPMs. If you use the High Availability port, specify the **internal** IP address (1.1.46.20) associated with the remote Primary CPM (obtained by running `show-internal-ip` on the remote chassis). If you use either or both of the management ports, specify the **external** IP address(es) associated with the port(s). You can specify all three addresses within a single `configure remote-box` command. Crossbeam recommends that you connect the CPMs using separate network broadcast domains for each pair of ports (the HA ports, the Management 1 ports, and the Management 2 ports). Connecting the ports directly can lead to scenarios that do not provide full redundancy.

Command:

```
CBS# configure remote-box 46 1.1.46.20 192.168.50.46 192.168.51.56  
CBS(conf-remote-box) # end
```

**NOTE:** The example `configure remote-box` command specifies only the IP addresses of the management 1 and 2 interfaces for one of the CPMs on chassis 2. If you have connected the management interfaces of the other CPM on Chassis 2, you can add the IP addresses for those interfaces.

### 1.1.3 Configuring the Synchronization Circuit

Configure a synchronization circuit between VAP Group fw1 on Chassis 1 and fw2 on Chassis 2 so that the two VAP Groups act as one Check Point Cluster with 6 members.

**NOTE:** This step must be performed **after** you configure the system-id, because the system-id affects the MAC selection and configuration of every circuit that gets created.

**NOTE:** When you configure a circuit, you can either assign a circuit-id number (1-4095) or let XOS assign the next available circuit-id number. This example illustrates the manual configuration method.

Enter these commands:

```
CBS# configure circuit sync circuit-id 100
CBS(conf-cct) # device-name sync
CBS(conf-cct) # internal
CBS(conf-cct) # vap-group fw1
CBS(conf-cct-vapgroup) # ip 193.0.0.11/24 increment-per-vap 193.0.0.13
CBS(conf-cct-vapgroup) # end
CBS#
CBS# configure interface gigabitethernet 4/2
CBS(conf-interface-gig) # logical sync
CBS(confi-gig-logical) # circuit sync
CBS(intf-logical) # end
CBS#
```

### 1.1.4 Configure the Traffic Circuits

Configure the two circuits that convey traffic to and from the fw1 VAP group.

**NOTE:** When you configure a circuit, you can either assign a circuit-id number (1-4095) or let XOS assign the next available circuit-id number. This example illustrates the manual configuration method.

```
CBS# configure circuit gig_18 circuit-id 108 device-name gig.18
CBS# configure circuit gig_18 vap-group fw1
CBS(conf-cct-vapgroup) # ip 172.16.20.2/24
CBS(conf-cct-vapgroup-ip) # enable
CBS(conf-cct-vapgroup-ip) # end
CBS# configure circuit gig_28 circuit-id 208 device-name gig.28
CBS# configure circuit gig_28 vap-group fw1
CBS(conf-cct-vapgroup) # ip 10.0.0.102/24
CBS(conf-cct-vapgroup-ip) # enable
CBS(conf-cct-vapgroup-ip) # end
CBS#
```

### 1.1.5 Configure the Traffic Interfaces

Configure the interfaces through which traffic flows to and from the fw1 VAP group.

```
CBS# configure interface gigabitethernet 1/8
CBS(conf-intf-gig) # logical gig18
CBS(intf-gig-logical) # circuit gig_18
CBS(intr-gig-log-cct) # end
CBS#
CBS# configure interface gigabitethernet 2/8
CBS(conf-intf-gig) # logical gig28
CBS(intf-gig-logical) # circuit gig_28
CBS(intr-gig-log-cct) # end
CBS#
```

## 1.2 Configuring Failover Group 1 (Chassis 1)

### 1.2.1 VRRP Failover Group

Create the failover group by assigning it a name (`failover1`) and a failover group ID. The failover group ID is different than the system identifier, configured earlier. The ID must be unique on this chassis, and must be the same on its counterpart failover group on the remote chassis (Chassis 2).

The `failover1` group acts as the master group on Chassis 1. The counterpart failover group on Chassis 2 is also called `failover1` and has the same group ID (1). The two groups have different `priority` values.

Command:

```
CBS# configure vrrp failover-group failover1 failover-group-id 1
CBS(conf-vrrp-group) #
```

### 1.2.2 VRRP Priority

For proper operation, the VRRP `priority` value of the two associated failover-groups must be different on Chassis 1 and Chassis 2; during normal operations, the failover group with the higher priority is the master. Certain events such as an interface failure or a change in VAP group member count can be configured to decrement the VRRP `priority` of the failover-group. Failover occurs when the VRRP `priority` value of one failover group drops below the priority of the failover group on the other chassis. VRRP `priority` values range from 1 to 255, and the default is 100.

Command:

```
CBS(conf-vrrp-group) # priority 250
CBS(conf-vrrp-group) # exit
```

### 1.2.3 Virtual Router on each Traffic Circuit

Create a virtual router on each traffic circuit that is attached to the `fw1` VAP group. A virtual router is assigned a virtual IP addresses that is used to configure VRRP and, optionally, next hop health check. This section describes the configuration of two virtual routers, one for each of the two circuits (`gig_18` and `gig_28`) that are associated with the `fw1` vap-group.

## Configuring the virtual-router for the `gig_18` circuit

1. To create the virtual router for the first circuit, enter this command:

```
CBS(conf-vrrp-group) # virtual-router vrrp-id 10 circuit gig_18
CBS(conf-vrrp-failover-vr) #
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr) # priority-delta 2
CBS(conf-vrrp-failover-vr) #
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 1 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 2, so the failover group on Chassis 2 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router

```
CBS(conf-vrrp-failover-vr) # mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr) #
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router

```
CBS(conf-vrrp-failover-vr) # vap-group fw1
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** Before you map the virtual router to the VAP group, the circuit must have been mapped to the VAP group.

Mapping the VAP group to the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 172.16.20.1/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

## Configuring the virtual-router for the gig\_28 circuit

1. To create the virtual router for the second circuit, enter this command:

```
CBS# configure vrrp failover-group failover1 failover-group-id 1
CBS(conf-vrrp-group) # virtual-router vrrp-id 11 circuit gig_28
CBS(conf-vrrp-failover-vr) #
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr) # priority-delta 2
CBS(conf-vrrp-failover-vr) #
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 1 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 2, so the failover group on Chassis 2 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router

```
CBS(conf-vrrp-failover-vr) # mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr) #
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router

```
CBS(conf-vrrp-failover-vr) # vap-group fw1
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** The circuit must already be mapped to the VAP group.

Specifying the VAP group of the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 10.0.0.101/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

## 1.2.4 Enable VRRP on the VAP Group

VRRP monitors the **fw1** VAP group for failure of individual VAPs. By setting the `active-vap-threshold` and the `priority-delta`, individual VAP failures can decrement VRRP priority and cause a failover. Failover occurs when the `priority-delta` value decrements the `priority` value of the master failover group below the `priority` value of the associated failover group on the remote chassis. In our configuration, the `priority` value for the master failover group is 250 and the `priority` value for the backup group is 249. A `priority-delta` of 2 is used for each of the VAP groups.

To configure the **fw1** VAP group for failover:

1. Enable VRRP on the **fw1** VAP group.

Command:

```
CBS# configure vrrp vap-group fw1
CBS(conf-vrrp-vap-group) #
```

2. Assign the VRRP enabled **fw1** VAP group to a failover group list (required).

Command:

```
CBS(conf-vrrp-vap-group) # failover-group-list failover1
CBS(conf-vrrp-vap-group) #
```

3. Optionally, set the hold down timer.

Configure the `hold-down-timer` to 120, which forces the VAP group to wait for two minutes while the application fully boots. This wait prevents the failover group from becoming VRRP master before the application is fully active.

Command:

```
CBS(conf-vrrp-vap-group) # hold-down-timer 120
CBS(conf-vrrp-vap-group) #
```

4. Optionally, set the active VAP threshold and return to the main CLI context.

The `active-vap-threshold` monitors the VAPs in the VAP group. If the number of active VAPs drops below the threshold, the group's `priority` value is decremented by the `priority-delta`. Failover occurs when the `priority-delta` decrements the `priority` value of the master failover group below the `priority` value of the backup group.

Command:

```
CBS(conf-vrrp vap-group) # active-vap-threshold 3
CBS(conf-vrrp vap-group) # end
CBS#
```

5. Set the `priority-delta` value for the VAP group (optional).

Assign a `priority-delta` to the VAP group. VRRP decrements the priority of the failover group whenever the number of active VAPs falls below the `active-vap-threshold`.

The `priority-delta` value can be any number between 1 and 255 and the default value is 1. When the VAP returns to the Active state, the `priority-delta` value is added back to the `priority` value.

Command:

```
CBS(conf-vrrp-vap-group) # priority-delta 2
CBS(conf-vrrp-vap-group) #
```

## 1.3 Optionally Configuring OSPF

If your network is configured to use the Open Shortest Path First (OSPF) protocol, you can incorporate OSPF into your VRRP configuration.

**NOTE:** To configure OSPF, you must first install the Crossbeam Routing Software (RSW).

When a failover occurs from one failover group to another, you want traffic to be rerouted from the failed group to the one that is now active. To ensure that this happens, you can increase the `ospf-cost-increment` value associated with the circuit on the first failover group. The new value is propagated to all local routers, increasing the OSPF cost of the circuit so that it is no longer part of the preferred route. When the original failover group resumes master status, the OSPF cost is readjusted to the originally configured value.

**NOTE:** Configure the `ospf-cost-increment` only on the master failover group.

To include OSPF cost in the configuration, perform these steps:

1. Configure these parameters on the master failover group (failover1):

```
CBS# configure vrrp failover-group failover1
CBS(conf-vrrp-group)# ospf-cost-increment circuit gig_18 5
CBS(conf-vrrp-group)# ospf-cost-increment circuit gig_28 5
```

2. On the VAP group that is associated with the master failover group, start the ospf routing protocol.

```
CBS# configure routing-protocol ospf vap-group fw1 start
```

## 1.4 Preemption

Typically, after a failover has occurred from one failover group to another, you want the failover group that is now master to remain in that role, even after the problem that caused the failover has been resolved.

By default, preemption is turned off, and Crossbeam recommends that you do not configure preemption for failover groups.

However, if you want the original failover group to resume the role of master, you can turn preemption on using this command:

```
CBS(conf-vrrp-group)# preemption
```

## 1.5 Management Circuit

If you intend to run Check Point software on the X-Series chassis and you intend to use a Check Point management station, do not configure the X-Series chassis management circuits as part of any failover group. If you do, the Check Point management station cannot access the management circuit.

## 1.6 Configuring Failover Group 2 (Chassis 1)

### 1.6.1 VRRP Failover Group

Create the failover group by assigning it a name (`failover2`) and a failover group ID. The failover group ID is different than the system identifier, configured earlier. The ID must be unique on this chassis, and must be the same on its counterpart failover group on the remote chassis (Chassis 2).

The failover2 group on Chassis 1 acts as the backup group to the master failover2 group on Chassis 2. Both groups have the same group ID (2). The two groups have different `priority` values.

Command:

```
CBS# configure vrrp failover-group failover2 failover-group-id 2
CBS(conf-vrrp-group)#
```

## 1.6.2 VRRP Priority

For proper operation, the VRRP `priority` value of the two associated failover-groups must be different on Chassis 1 and Chassis 2; during normal operations, the failover group with the higher priority is the master. Certain events such as an interface failure or a change in VAP group member count can be configured to decrement the VRRP priority of the failover-group. Failover occurs when the VRRP `priority` value of one failover group drops below the priority of the failover group on the other chassis. VRRP `priority` values range from 1 to 255, and the default is 100.

Command:

```
CBS(conf-vrrp-group)# priority 249
```

## 1.6.3 Virtual Router on each Traffic Circuit

Create a virtual router on each traffic circuit that is attached to the **fw1** VAP group. A virtual router is assigned a virtual IP addresses that is used to configure VRRP and, optionally, next hop health check. This section describes the configuration of two virtual routers, one for each of the two circuits (`gig_18` and `gig_28`) that are associated with the **fw1** VAP group.

## Configuring the virtual-router for the `gig_18` circuit

1. To create the virtual router for the first circuit, enter this command:

```
CBS(conf-vrrp-group)# virtual-router vrrp-id 20 circuit gig_18
CBS(conf-vrrp-failover-vr)#
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr)# priority-delta 2
CBS(conf-vrrp-failover-vr)#
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 2 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 1, so the failover group on Chassis 1 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router

```
CBS(conf-vrrp-failover-vr)# mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr)#
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router

```
CBS(conf-vrrp-failover-vr)# vap-group fw1
CBS(conf-vrrp-vr-vapgroup)#
```

**NOTE:** Before you map the virtual router to the VAP group, the circuit must have been mapped to the VAP group.

Mapping the VAP group to the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 172.0.10.1/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

## Configuring the virtual-router for the gig\_28 circuit

1. To create the virtual router for the second circuit, enter this command:

```
CBS# configure vrrp failover-group failover2 failover-group-id 2
CBS(conf-vrrp-group) # virtual-router vrrp-id 21 circuit gig_28
CBS(conf-vrrp-failover-vr) #
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr) # priority-delta 2
CBS(conf-vrrp-failover-vr) #
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 1 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 2, so the failover group on Chassis 2 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router

```
CBS(conf-vrrp-failover-vr) # mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr) #
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router

```
CBS(conf-vrrp-failover-vr) # vap-group fw1
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** The circuit must already be mapped to the VAP group.

Specifying the VAP group of the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 10.30.10.101/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

### 1.6.4 Enable VRRP on the VAP Group

VRRP monitors the `fw1` VAP group for failure of individual VAPs. By setting the `active-vap-threshold` and the `priority-delta`, individual VAP failures can decrement VRRP priority and cause a comparison with the VRRP priority on the remote chassis. Failover occurs when the `priority-delta` value decrements the `priority` value of the master failover group below the `priority` value of the associated failover group on the remote chassis. In our configuration, the `priority` value for the master failover group is 250 and the `priority` value for the backup group is 249. A `priority-delta` of 2 is used for each of the VAP groups.

To configure the `fw1` VAP group for failover:

1. Enable VRRP on the **fw1** VAP group.

```
CBS# configure vrrp vap-group fw1  
CBS(conf-vrrp-vap-group) #
```

2. Assign the VRRP enabled **fw1** VAP group to a failover group list (required).

```
CBS(conf-vrrp-vap-group) # failover-group-list failover2  
CBS(conf-vrrp-vap-group) #
```

3. Optionally, set the hold down timer.

Configure the `hold-down-timer` to 120, which forces the VAP group to wait for two minutes while the application fully boots. This wait prevents the failover group from becoming VRRP master before the application is fully active.

```
CBS(conf-vrrp-vap-group) # hold-down-timer 120  
CBS(conf-vrrp-vap-group) #
```

4. Optionally, set the active VAP threshold and return to the main CLI context.

The `active-vap-threshold` monitors the number of active VAPs in the VAP group. If the number of active VAPs drops below the threshold, the failover group's `priority` value is decremented by the `priority-delta` (defined in [Step 5](#)) and a comparison is done between the priorities of the failover groups on the two chassis. Failover occurs when the `priority-delta` decrements the `priority` value of the master failover group below the `priority` value of the backup group.

```
CBS(conf-vrrp vap-group) # active-vap-threshold 3  
CBS(conf-vrrp vap-group) # end  
CBS#
```

5. Set the `priority-delta` value for the VAP group (optional).

Assign a `priority-delta` to the VAP group. VRRP decrements the priority of the failover group whenever the number of active VAPs falls below the `active-vap-threshold`.

The `priority-delta` value can be any number between 1 and 255 and the default value is 1. When the VAP returns to the Active state, the `priority-delta` value is added back to the `priority` value.

```
CBS(conf-vrrp-vap-group) # priority-delta 2  
CBS(conf-vrrp-vap-group) # exit
```

**NOTE:** Do not configure any OSPF cost increments for the circuits associated with `failover2` group on chassis 1. OSPF is configured only for master failover groups.

## 1.7 Preemption

Typically, after a failover has occurred from one failover group to another, you want the failover group that is now master to remain in that role, even after the problem that caused the failover has been resolved.

By default, preemption is turned off, and Crossbeam recommends that you do not configure preemption for failover groups.

However, if you want the original failover group to resume the role of master, you can turn preemption on using this command:

```
CBS(conf-vrrp-group) # preemption
```

## 1.8 Management Circuit

If you intend to run Check Point software on the X-Series chassis and you intend to use a Check Point management station, do not configure the X-Series chassis management circuits as part of any failover group. If you do, the Check Point management station cannot access the management circuit.

## Verifying Your Configuration

This section contains the output of several commands that show the configured status of Chassis 1. To check your progress throughout the configuration process, open a second CLI window, log in to the CPM and enter one of the commands.

### Output of show running-config on Chassis 1

```
CBS# show running-config
#
vap-group fw1 xslinux_v5
  vap-count 3
  max-load-count 3
  ap-list ap1 ap2 ap3 ap4 ap5 ap6 ap7 ap8 ap9 ap10
  load-balance-vap-list 1 2 3 4 5 6 7 8 9 10
  ip-flow-rule loadbalance1
    action load-balance
    activate
#
circuit gig_18 circuit-id 108
  device-name gig.18
  vap-group fw1
    ip 172.16.20.2/24 172.16.20.255
circuit gig_28 circuit-id 208
  device-name gig.28
  vap-group fw1
    ip 10.0.0.102/24 10.0.0.255
#
interface gigabitethernet 1/8
  logical gig18
    circuit gig_18
interface gigabitethernet 2/8
  logical gig28
    circuit gig_28
#
#
vrrp failover-group failover1 failover-group-id 1
  priority 250
  virtual-router vrrp-id 10 circuit gig_18
    priority-delta 2
    vap-group fw1
      virtual-ip 172.16.20.1/24 172.16.20.255
  virtual-router vrrp-id 11 circuit gig_28
    priority-delta 2
    vap-group fw1
      virtual-ip 10.0.0.101/24 10.0.0.255
vrrp failover-group failover2 failover-group-id 2
  priority 249
  virtual-router vrrp-id 20 circuit gig_18
    priority-delta 2
    vap-group fw1
      virtual-ip 172.0.10.1/24 172.16.20.255
  virtual-router vrrp-id 21 circuit gig_28
    priority-delta 2
    vap-group fw1
      virtual-ip 10.30.10.101/24 10.0.0.255
#
vrrp vap-group fw1
```

```

failover-group-list failover1 failover2
hold-down-timer 120
priority-delta 2
#
management gigabitethernet 13/1
ip-addr 192.168.50.45/24 192.168.50.255
enable
access-list 1 input
access-list 2 output
management gigabitethernet 13/2
ip-addr 192.168.51.55/24 192.168.51.255
enable
access-list 1 input
access-list 2 output
management gigabitethernet 14/1
ip-addr 192.168.50.65/24 192.168.50.255
enable
access-list 1 input
access-list 2 output
management gigabitethernet 14/2
ip-addr 192.168.51.75/24 192.168.51.255
enable
access-list 1 input
access-list 2 output

```

### Output of show vrrp on Chassis 1

```

CBS# show vrrp
Priority is Actual/Configured
FG-ID  Priority  Status  Preempt  Master Sys ID  Master Priority
1      250/250    Master  off      45             250
2      249/249    Backup off      46             250
(1 row)

```

## Output of show vrrp detail-status on Chassis 1

```
CBS# show vrrp detail-status
FG_ID  Status  Priority  Delta  Type  Component
  1    Master  250/250    2     vr   gig_18/10/5
  1    Master  250/250    2     vr   gig_28/10/5
  1    Master  250/250    2     vg   fw1
  2    Backup  249/249    2     vr   gig_18/10/0
  2    Backup  249/249    2     vr   gig_28/10/0
  2    Backup  249/249    2     vg   fw1
```

**NOTE:** The Component column shows 5 as the last digit only if you configured the OSPF cost increment as described in [Optionally Configuring OSPF](#) on page 58. If you did not configure the OSPF cost increment, these values would be 0 (zero).

## Output of show remote-box on Chassis 1

```
CBS# show remote-box
System ID      : 46
First IP Address : 1.1.46.20
Second IP Address : 192.168.50.46
Third IP Address  : 192.168.51.56
Fourth IP Address : 0.0.0.0
Fifth IP Address  : 0.0.0.0
Active IP Address : 1.1.46.20
(1 row)
```

## 2.0 Configuring Chassis 2

On Chassis 2, perform these tasks:

- [Configuring System-wide Parameters \(Chassis 2\)](#)
- [Configuring Failover Group 1 \(Chassis 2\)](#)
- [Configuring Failover Group 2 \(Chassis 2\)](#)

### 2.1 Configuring System-wide Parameters (Chassis 2)

#### 2.1.1 Local System Identifier

Configure the local system-identifier on Chassis 2.

**NOTE:** When configuring multiple chassis for high availability, a unique system ID must be assigned to each chassis. If both chassis are configured with the same ID, you run the risk of having identical MAC addresses on any given circuit. This configuration is **not** supported or recommended.

If your X-Series Platforms do not have unique system IDs assigned, use the following command to define a system ID. The valid range is from 1-255.

Command:

```
CBS# configure system-identifier 46
```

**NOTE:** After you configure a system-identifier, you must use the `reload all` command to activate the identifier.

#### 2.1.2 Remote System Identifier

**NOTE:** If you connect Chassis 1 to Chassis 2 using only the High Availability (HA) port, you do not need to configure the remote system identifier. The two chassis use the HA port to exchange system-id information. However, if you do this, the HA port connection represents a single point of failure. Crossbeam recommends that you connect the HA ports and at least one pair of management interfaces.

Configure the remote system ID and IP address using the `configure remote-box` command.

**NOTE:** The `remote-box` command requires that you have interconnected CPMs on the two chassis. The IP address that you specify depends on how you have interconnected the CPMs. If you use the High Availability port, specify the **internal** IP address (1.1.46.20) associated with the remote Primary CPM (obtained by running `show-internal-ip` on the remote chassis). If you use either or both of the management ports, specify the **external** IP address(es) associated with the port(s). You can specify all three addresses within a single `configure remote-box` command. Crossbeam recommends that you connect the CPMs using separate network broadcast domains for each pair of ports (the HA ports, the Management 1 ports, and the Management 2 ports). Connecting the ports directly can lead to scenarios that do not provide full redundancy.

Command:

```
CBS# configure remote-box 45 1.1.45.20 192.168.50.45 192.168.51.55  
CBS(conf-remote-box) # end  
CBS#
```

**NOTE:** The example `configure remote-box` command specifies only the IP addresses of the management 1 and 2 interfaces for one of the CPMs on chassis 1. If you have connected the management interfaces of the other CPM on Chassis 2, you can add the IP addresses for those interfaces.

### 2.1.3 Configuring the Synchronization Circuit

Configure a synchronization circuit between VAP Group fw2 on Chassis 2 and fw1 on Chassis 1 so that the two VAP Groups act as one Check Point Cluster with 6 members.

**NOTE:** This step must be performed **after** you configure the system-id, because the system-id affects the MAC selection and configuration of every circuit that gets created.

**NOTE:** When you configure a circuit, you can either assign a circuit-id number (1-4095) or let XOS assign the next available circuit-id number. This example illustrates the manual configuration method.

Enter these commands:

```
CBS# configure circuit sync circuit-id 100
CBS(conf-cct) # device-name sync
CBS(conf-cct) # internal
CBS(conf-cct) # vap-group fw2
CBS(conf-cct-vapgroup) # ip 193.0.0.14/24 increment-per-vap 193.0.0.16
CBS(conf-cct-vapgroup) # end
CBS#
CBS# configure interface gigabitethernet 4/2
CBS(conf-interface-gig) # logical sync
CBS(confi-gig-logical) # circuit sync
CBS(intf-logical) # end
CBS#
```

### 2.1.4 Configure the Traffic Circuits

Configure the two circuits that convey traffic to and from the fw1 VAP group.

**NOTE:** When you configure a circuit, you can either assign a circuit-id number (1-4095) or let XOS assign the next available circuit-id number. This example illustrates the manual configuration method.

```
CBS# configure circuit gig_17 circuit-id 107 device-name gig.17
CBS# configure circuit gig_17 vap-group fw2
CBS(conf-cct-vapgroup) # ip 172.16.20.3/24
CBS(conf-cct-vapgroup-ip) # enable
CBS(conf-cct-vapgroup-ip) # end
CBS# configure circuit gig_27 circuit-id 207 device-name gig.27
CBS# configure circuit gig_27 vap-group fw2
CBS(conf-cct-vapgroup) # ip 10.0.0.103/24
CBS(conf-cct-vapgroup-ip) # enable
CBS(conf-cct-vapgroup-ip) # end
CBS#
```

### 2.1.5 Configure the Traffic Interfaces

Configure the interfaces through which traffic flows to and from the fw1 VAP group.

```
CBS# configure interface gigabitethernet 1/7
CBS(conf-intf-gig) # logical gig17
CBS(intf-gig-logical) # circuit gig_17
CBS(intf-gig-logical-cct) # end
CBS#
CBS# configure interface gigabitethernet 2/7
CBS(conf-intf-gig) # logical gig27
CBS(intf-gig-logical) # circuit gig_27
CBS(intf-gig-logical-cct) # end
CBS#
```

## 2.2 Configuring Failover Group 1 (Chassis 2)

### 2.2.1 VRRP Failover Group

Create the failover group by assigning it a name (`failover1`) and a failover group ID. The failover group ID is different than the system identifier, configured earlier. The ID must be unique on this chassis, and must be the same on its counterpart failover group on the remote chassis (Chassis 1).

The `failover1` group on Chassis 2 acts as the backup group to the master `failover1` group on Chassis 1. Both groups have the same group ID (1). The two groups have different `priority` values.

Command:

```
CBS# configure vrrp failover-group failover1 failover-group-id 1
CBS(conf-vrrp-group)#
```

### 2.2.2 VRRP Priority

For proper operation, the VRRP `priority` value of the two associated failover-groups must be different on Chassis 1 and Chassis 2; during normal operations, the failover group with the higher priority is the master. Certain events such as an interface failure or a change in VAP group member count can be configured to decrement the VRRP priority of the failover-group. Failover occurs when the VRRP `priority` value of one failover group drops below the priority of the failover group on the other chassis. VRRP `priority` values range from 1 to 255, and the default is 100.

Command:

```
CBS(conf-vrrp-group)# priority 249
CBS(conf-vrrp-group)# exit
```

### 2.2.3 Virtual Router on each Traffic Circuit

Create a virtual router on each traffic circuit that is attached to the `fw2` VAP group. A virtual router is assigned a virtual IP addresses that is used to configure VRRP and, optionally, next hop health check. This section describes the configuration of two virtual routers, one for each of the two circuits (`gig_17` and `gig_27`) that are associated with the `fw2` vap-group.

## Configuring the virtual-router for the `gig_17` circuit

1. To create the virtual router for the first circuit, enter this command:

```
CBS(conf-vrrp-group)# virtual-router vrrp-id 10 circuit gig_17
CBS(conf-vrrp-failover-vr)#
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr)# priority-delta 2
CBS(conf-vrrp-failover-vr)#
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 1 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 2, so the failover group on Chassis 2 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router.

```
CBS(conf-vrrp-failover-vr)# mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr)#
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router.

```
CBS(conf-vrrp-failover-vr) # vap-group fw2
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** Before you map the virtual router to the VAP group, the circuit must have been mapped to the VAP group.

Mapping the VAP group to the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 172.16.20.1/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

## Configuring the virtual-router for the gig\_27 circuit

1. To create the virtual router for the second circuit, enter this command:

```
CBS# configure vrrp failover-group failover1 failover-group-id 1
CBS(conf-vrrp-group) # virtual-router vrrp-id 11 circuit gig_27
CBS(conf-vrrp-failover-vr) #
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr) # priority-delta 2
CBS(conf-vrrp-failover-vr) #
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 1 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 2, so the failover group on Chassis 2 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router,

```
CBS(conf-vrrp-failover-vr) # mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr) #
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router,

```
CBS(conf-vrrp-failover-vr) # vap-group fw2
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** The circuit must already be mapped to the VAP group.

Specifying the VAP group of the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS (conf-vrrp-vr-vapgroup) # virtual-ip 10.0.0.101/24
CBS (conf-vrrp-vr-vapgroup) # end
CBS#
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

## 2.2.4 Enable VRRP on the VAP Group

VRRP monitors the **fw2** VAP group for failure of individual VAPs. By setting the `active-vap-threshold` and the `priority-delta`, individual VAP failures can decrement VRRP priority and cause a failover. Failover occurs when the `priority-delta` value decrements the `priority` value of the master failover group below the `priority` value of the associated failover group on the remote chassis. In our configuration, the `priority` value for the master failover group is 250 and the `priority` value for the backup group is 249. A `priority-delta` of 2 is used for each of the VAP groups.

To configure the **fw2** VAP group for failover:

1. Enable VRRP on the **fw2** VAP group.

Command:

```
CBS# configure vrrp vap-group fw2
CBS (conf-vrrp-vap-group) #
```

2. Assign the VRRP enabled **fw2** VAP group to a failover group list (required).

Command:

```
CBS (conf-vrrp-vap-group) # failover-group-list failover1
CBS (conf-vrrp-vap-group) #
```

3. Optionally, set the hold down timer.

Configure the `hold-down-timer` to 120, which forces the VAP group to wait for two minutes while the application fully boots. This wait prevents the failover group from becoming VRRP master before the application is fully active.

Command:

```
CBS (conf-vrrp-vap-group) # hold-down-timer 120
CBS (conf-vrrp-vap-group) #
```

4. Optionally, set the active VAP threshold and return to the main CLI context.

The `active-vap-threshold` monitors the VAPs in the VAP group. If the number of active VAPs drops below the threshold, the group's `priority` value is decremented by the `priority-delta`. Failover occurs when the `priority-delta` decrements the `priority` value of the master failover group below the `priority` value of the backup group.

Command:

```
CBS (conf-vrrp vap-group) # active-vap-threshold 3
CBS (conf-vrrp vap-group) # end
CBS#
```

5. Set the `priority-delta` value for the VAP group (optional).

Assign a `priority-delta` to the VAP group. VRRP decrements the priority of the failover group whenever the number of active VAPs falls below the `active-vap-threshold`.

The `priority-delta` value can be any number between 1 and 255 and the default value is 1. When the VAP returns to the Active state, the `priority-delta` value is added back to the `priority` value.

Command:

```
CBS (conf-vrrp-vap-group) # priority-delta 2
CBS (conf-vrrp-vap-group) # exit
```

**NOTE:** Do not configure any OSPF cost increments for the circuits associated with failover1 group on chassis 2. OSPF is configured only for master failover groups.

## 2.3 Preemption

Typically, after a failover has occurred from one failover group to another, you want the failover group that is now master to remain in that role, even after the problem that caused the failover has been resolved.

By default, preemption is turned off, and Crossbeam recommends that you do not configure preemption for failover groups.

However, if you want the original failover group to resume the role of master, you can turn preemption on using this command:

```
CBS(conf-vrrp-group) # preemption
```

## 2.4 Management Circuit

If you intend to run Check Point software on the X-Series chassis and you intend to use a Check Point management station, do not configure the X-Series chassis management circuits as part of any failover group. If you do, the Check Point management station cannot access the management circuit.

## 2.5 Configuring Failover Group 2 (Chassis 2)

### 2.5.1 VRRP Failover Group

Create the failover group by assigning it a name (*failover2*) and a failover group ID. The failover group ID is different than the system identifier, configured earlier. The ID must be unique on this chassis, and must be the same on its counterpart failover group on the remote chassis (Chassis 1).

The failover2 group acts as the master group on Chassis 2. The counterpart failover group on Chassis 1 is also called failover2 and has the same group ID (2). The two groups have different *priority* values.

Command:

```
CBS# configure vrrp failover-group failover2 failover-group-id 2  
CBS(conf-vrrp-group) #
```

### 2.5.2 VRRP Priority

For proper operation, the VRRP *priority* value of the two associated failover-groups must be different on Chassis 1 and Chassis 2; during normal operations, the failover group with the higher priority is the master. Certain events such as an interface failure or a change in VAP group member count can be configured to decrement the VRRP priority of the failover-group. Failover occurs when the VRRP *priority* value of one failover group drops below the priority of the failover group on the other chassis. VRRP *priority* values range from 1 to 255, and the default is 100.

Command:

```
CBS(conf-vrrp-group) # priority 250  
CBS(conf-vrrp-group) #
```

### 2.5.3 Virtual Router on each Traffic Circuit

Create a virtual router on each traffic circuit that is attached to the **fw2** VAP group. A virtual router is assigned a virtual IP addresses that is used to configure VRRP and, optionally, next hop health check. This section describes the configuration of two virtual routers, one for each of the two circuits (*gig\_17* and *gig\_27*) that are associated with the **fw2** vap-group.

## Configuring the virtual-router for the gig\_17 circuit

1. To create the virtual router for the first circuit, enter this command:

```
CBS(conf-vrrp-group)# virtual-router vrrp-id 20 circuit gig_17  
CBS(conf-vrrp-failover-vr)#
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr)# priority-delta 2  
CBS(conf-vrrp-failover-vr)#
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of the failover groups on the two chassis. In this example, the priority value on Chassis 2 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 1, so the failover group on Chassis 1 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router,

```
CBS(conf-vrrp-failover-vr)# mac-usage vrrp-mac  
CBS(conf-vrrp-failover-vr)#
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router,

```
CBS(conf-vrrp-failover-vr)# vap-group fw2  
CBS(conf-vrrp-vr-vapgroup)#
```

**NOTE:** Before you map the virtual router to the VAP group, the circuit must have been mapped to the VAP group.

Mapping the VAP group to the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup)# virtual-ip 172.0.10.1/24  
CBS(conf-vrrp-vr-vapgroup)# end  
CBS#
```

## Configuring the virtual-router for the gig\_27 circuit

1. To create the virtual router for the second circuit, enter this command:

```
CBS(conf-vrrp-group)# virtual-router vrrp-id 21 circuit gig_27  
CBS(conf-vrrp-failover-vr)#
```

**NOTE:** Any circuit that is defined as part of a virtual router is automatically monitored for failure, and, when one occurs, the failover group's priority is decremented by the `priority-delta` value.

2. Assign a `priority-delta` value to the virtual router.

```
CBS(conf-vrrp-failover-vr)# priority-delta 2  
CBS(conf-vrrp-failover-vr)#
```

When a virtual router fails, the associated failover group's `priority` value is decremented by the `priority-delta` value (2) to a new priority value and a comparison is done between the priorities of

the failover groups on the two chassis. In this example, the priority value on Chassis 1 becomes 248, which is less than the `priority` value of the associated failover group on Chassis 2, so the failover group on Chassis 2 becomes the master. The `priority-delta` value is added back to the priority when the VR recovers.

3. Specify MAC usage on the VRRP Virtual Router.

```
CBS(conf-vrrp-failover-vr) # mac-usage vrrp-mac
CBS(conf-vrrp-failover-vr) #
```

When you specify `vrrp-mac` as the MAC usage parameter, instead of using the same MAC address on both the circuit and virtual IP address, XOS automatically generates a unique `vrrp-mac` (for example, `00:00:5e:00:01:10` where the last two digits is the VRRP ID of the Virtual Router). In the event of a failover, the MAC address moves with the virtual IP address to the new chassis. By keeping the MAC address consistent, it reduces the time required for upstream routers to converge routing tables.

4. Specify the VAP Group of the Virtual Router.

```
CBS(conf-vrrp-failover-vr) # vap-group fw2
CBS(conf-vrrp-vr-vapgroup) #
```

**NOTE:** The circuit must already be mapped to the VAP group.

Specifying the VAP group of the virtual router allows you to assign a virtual IP address to the VAP group.

5. Assign a virtual IP Address to the Virtual Router, and return to the main CLI context.

```
CBS(conf-vrrp-vr-vapgroup) # virtual-ip 10.30.10.101/24
CBS(conf-vrrp-vr-vapgroup) # end
CBS#
```

## 2.5.4 Enable VRRP on the VAP Group

VRRP monitors the `fw2` VAP group for failure of individual VAPs. By setting the `active-vap-threshold` and the `priority-delta`, individual VAP failures can decrement VRRP priority and cause a comparison with the VRRP priority on the remote chassis. Failover occurs when the `priority-delta` value decrements the `priority` value of the master failover group below the `priority` value of the associated failover group on the remote chassis. In our configuration, the `priority` value for the master failover group is 250 and the `priority` value for the backup group is 249. A `priority-delta` of 2 is used for each of the VAP groups.

To configure the `fw2` VAP group for failover:

1. Enable VRRP on the `fw2` VAP group.

```
CBS# configure vrrp vap-group fw2
CBS(conf-vrrp-vap-group) #
```

2. Assign the VRRP enabled `fw2` VAP group to a failover group list (required).

```
CBS(conf-vrrp-vap-group) # failover-group-list failover2
CBS(conf-vrrp-vap-group) #
```

3. Optionally, set the hold down timer.

Configure the `hold-down-timer` to 120, which forces the VAP group to wait for two minutes while the application fully boots. This wait prevents the failover group from becoming VRRP master before the application is fully active.

```
CBS(conf-vrrp-vap-group) # hold-down-timer 120
CBS(conf-vrrp-vap-group) #
```

4. Optionally, set the active VAP threshold and return to the main CLI context.

The `active-vap-threshold` monitors the number of active VAPs in the VAP group. If the number of active VAPs drops below the threshold, the failover group's `priority` value is decremented by the `priority-delta` (defined in [Step 5](#)) and a comparison is done between the priorities of the failover

groups on the two chassis. Failover occurs when the `priority-delta` decrements the `priority` value of the master failover group below the `priority` value of the backup group.

```
CBS(conf-vrrp vap-group) # active-vap-threshold 3
CBS(conf-vrrp vap-group) # end
CBS#
```

5. Set the `priority-delta` value for the VAP group (optional).

Assign a `priority-delta` to the VAP group. VRRP decrements the priority of the failover group whenever the number of active VAPs falls below the `active-vap-threshold`.

The `priority-delta` value can be any number between 1 and 255 and the default value is 1. When the VAP returns to the Active state, the `priority-delta` value is added back to the `priority` value.

```
CBS(conf-vrrp-vap-group) # priority-delta 2
CBS(conf-vrrp-vap-group) # exit
```

## 2.6 Optionally Configuring OSPF

If your network is configured to use the Open Shortest Path First (OSPF) protocol, you can incorporate OSPF into your VRRP configuration.

**NOTE:** To configure OSPF, you must first install the Crossbeam Routing Software (RSW).

When a failover occurs from one failover group to another, you want traffic to be rerouted from the failed group to the one that is now active. To ensure that this happens, you can increase the `ospf-cost-increment` value associated with the circuit on the first failover group. The new value is propagated to all local routers, increasing the OSPF cost of the circuit so that it is no longer part of the preferred route. When the original failover group resumes master status, the OSPF cost is readjusted to the originally configured value.

**NOTE:** Configure the `ospf-cost-increment` only on the master failover group.

To include OSPF cost in the configuration, perform these steps:

1. Configure these parameters on the master failover group (failover1):

```
CBS# configure vrrp failover-group failover1
CBS(conf-vrrp-group) # ospf-cost-increment circuit gig_18 5
CBS(conf-vrrp-group) # ospf-cost-increment circuit gig_28 5
CBS(conf-vrrp-group) # end
CBS#
```

2. On the VAP group that is associated with the master failover group, start the ospf routing protocol.

```
CBS# configure routing-protocol ospf vap-group fw1 start
```

## 2.7 Preemption

Typically, after a failover has occurred from one failover group to another, you want the failover group that is now master to remain in that role, even after the problem that caused the failover has been resolved.

By default, preemption is turned off, and Crossbeam recommends that you do not configure preemption for failover groups.

However, if you want the original failover group to resume the role of master, you can turn preemption on using this command:

```
CBS(conf-vrrp-group) # preemption
```

## 2.8 Management Circuit

If you intend to run Check Point software on the X-Series chassis and you intend to use a Check Point management station, do not configure the X-Series chassis management circuits as part of any failover group. If you do, the Check Point management station cannot access the management circuit.

### Verifying Your Configuration

This section contains the output of several commands that show the configured status of Chassis 2. To check your progress throughout the configuration process, open a second CLI window, log in to the CPM and enter one of the commands.

#### Output of show running-config on Chassis 2

```
CBS# show running-config
#
vap-group fw2 xslinux_v5
  vap-count 3
  max-load-count 3
  ap-list ap1 ap2 ap3 ap4 ap5 ap6 ap7 ap8 ap9 ap10
  load-balance-vap-list 1 2 3 4 5 6 7 8 9 10
  ip-flow-rule loadbalance1
    action load-balance
    activate
#
circuit gig_17 circuit-id 107
  device-name gig.17
  vap-group fw2
    ip 172.16.20.3/24 172.16.20.255
circuit gig_27 circuit-id 207
  device-name gig.27
  vap-group fw2
    ip 10.0.0.103/24 10.0.0.255
#
interface gigabitethernet 1/7
  logical gig17
    circuit gig_17
interface gigabitethernet 2/7
  logical gig27
    circuit gig_27
#
vrrp failover-group failover1 failover-group-id 1
  priority 249
  virtual-router vrrp-id 10 circuit gig_17
    priority-delta 2
    vap-group fw2
      virtual-ip 172.16.20.1/24 172.16.20.255
  virtual-router vrrp-id 11 circuit gig_27
    priority-delta 2
    vap-group fw2
      virtual-ip 10.0.0.101/24 10.0.0.255
```

```

vrrp failover-group failover2 failover-group-id 2
  priority 250
  virtual-router vrrp-id 20 circuit gig_17
    priority-delta 2
    vap-group fw2
      virtual-ip 172.0.10.1/24 172.16.20.255
  virtual-router vrrp-id 21 circuit gig_27
    priority-delta 2
    vap-group fw2
      virtual-ip 10.30.10.101/24 10.0.0.255
#
vrrp vap-group fw2
  failover-group-list failover1 failover2
  hold-down-timer 120
  priority-delta 2
#
management gigabitethernet 13/1
  ip-addr 192.168.50.45/24 192.168.50.255
  enable
  access-list 1 input
  access-list 2 output
management gigabitethernet 13/2
  ip-addr 192.168.51.55/24 192.168.51.255
  enable
  access-list 1 input
  access-list 2 output
management gigabitethernet 14/1
  ip-addr 192.168.50.65/24 192.168.50.255
  enable
  access-list 1 input
  access-list 2 output
management gigabitethernet 14/2
  ip-addr 192.168.51.75/24 192.168.51.255
  enable
  access-list 1 input
  access-list 2 output

```

## Output of show vrrp on Chassis 2

```

CBS# show vrrp
Priority is Actual/Configured
FG-ID  Priority  Status  Preempt  Master Sys ID  Master Priority
1      249/249    Backup  off      45             250
2      250/250    Master  off      46             250
(1 row)

```

## Output of show vrrp detail-status on Chassis 2

```
CBS# show vrrp detail-status
FG_ID  Status  Priority  Delta  Type  Component
  1    Backup  249/249    2     vr   gig_17/20/0
  1    Backup  249/249    2     vr   gig_27/21/0
  1    Backup  249/249    2     vg   fw2
  2    Master  250/250    2     vr   gig_17/20/5
  2    Master  250/250    2     vr   gig_27/21/5
  2    Master  250/250    2     vg   fw2
```

**NOTE:** The Component column shows 5 as the last digit only if you configured the OSPF cost increment as described in [Optionally Configuring OSPF](#) on page 43. If you did not configure the OSPF cost increment, these values would be 0 (zero).

## Output of show remote-box on Chassis 1

```
CBS# show remote-box
System ID      : 46
First IP Address : 1.1.46.20
Second IP Address : 192.168.50.46
Third IP Address  : 192.168.51.56
Fourth IP Address : 0.0.0.0
Fifth IP Address  : 0.0.0.0
Active IP Address : 1.1.46.20
(1 row)
```



# Appendix A

## Basic Chassis Configuration

This appendix describes the basic configuration of the two chassis. If you have already set up your chassis, you can skip this section and start with [Active-Standby Configuration](#) on page 18 or [Active-Active Configuration](#) on page 36.

### Assign Hostnames

Assign a hostname as follows:

#### Chassis 1

Command:

```
CBS # configure hostname <your hostname goes here> cp1
CBS # configure hostname <your hostname goes here> cp2
```

**NOTE:** If you do not specify cp1 or cp2, the hostname is applied to both CPMs.

#### Chassis 2

Command:

```
CBS # configure hostname <your hostname goes here> cp1
CBS # configure hostname <your hostname goes here> cp2
```

**NOTE:** If you do not specify cp1 or cp2, the hostname is applied to both CPMs.

### Assign a Domain Name

#### Chassis 1

Command:

```
CBS # configure ip domainname <your domain name goes here>
```

#### Chassis 2

Command:

```
CBS # configure ip domainname <your domain name goes here>
```



---

# Glossary

The following terms are used throughout the X-Series system documentation set.

## **3DES**

Triple Data Encryption Standard. Provides a stronger form of DES encryption where the algorithm is applied three times in order to encrypt data.

## **ACL**

Access Control List. Provides packet filtering through the permission or denial of packets based on certain IP criteria, such as IP address, port, or protocol.

## **APM**

Application Processor Module. The XOS Application Processor system module that provides application processing, status monitoring, and standard and application specific logging. The APM contains one or more CPUs to host applications and network services while processing packets belonging to individual flows.

## **ARP**

Address Resolution Protocol. An Internet protocol used to map an IP address to a MAC address.

## **BOTW**

Bump-on-the-Wire. A device with two or more interfaces that are transparent to the adjacent Layer 3 devices.

## **cbsflowagentd**

Flow Agent daemon that collects statistics and runs on each VAP.

## **cbsflowcalcd**

Flow Calculator daemon that runs the flow scheduling chow file and executes on the CPM.

## **circuit**

An abstract object representing a logical network interface (network service access point). A circuit can be mapped to either single or multiple logical lines. Attributes of a circuit include: a set of physical line or channel pairs, a Layer 2 encapsulation type, a Layer 2 address, and an IP address (optional).

## **CLI**

Command Line Interface.

## **CM**

Configuration manager/monitor.

## **core-intf**

An interface which is attached to the core-facing networks.

## **CPM**

Control Processor Module. The XOS system module that coordinates the actions of all other modules, enables management access to the platform, and supports access to a local disk containing configuration files and databases necessary to execute the applications which reside on the platform.

## **DES**

Data Encryption Standard. A popular algorithm for encrypting data. It is a product cipher that operates on 64-bit blocks of data, using a 56-bit key.

## **device**

OS concept representing either a physical or logical I/O port connected to the APM.

## **domain**

A set of interconnected IP networks belonging to a unique address space. A domain is uniquely identified within the X-Series system by a 8-bit domain ID. IP flows must be unique within a given domain.

## **DSA**

Digital Signature Algorithm.

## **ECC**

Error Checking and Correcting. A collection of methods to detect errors in transmitted or stored data and a means to correct them.

## **edge interface**

An interface that is attached to edge-facing networks (typically where subscribers are located).

## **edge server**

A server that is physically located close to its end users designed to deliver faster, higher quality transmissions, typically in a local commercial ISP facility. The number of edge servers in a region depends on the number of users in the locale.

## **Element Management System**

Graphical user interface for X-Series systems accessed via the web server.

## **Error Checking and Correcting (ECC)**

A collection of methods for detecting and correcting errors in transmitted or stored data.

## **FCAPS**

Faults, Configuration, Accounting, Performance, and Security. The general requirements of a network management system as defined by the International Organization for Standardization.

## **FIB**

Forwarding Information Base. A set of IP data structures replacing a route table in Linux.

## **firewall**

A set of software tools that protects a company's internal network from unwanted entry by unauthorized external users. The firewall works in conjunction with a router program to filter incoming network packets and reject those of unknown origin.

## **flow**

Specific stream of data traveling between two endpoints across a network. Specified by source IP, destination IP, source port, destination port and IP protocol type.

## **flow rule**

A filter rule specifying how a packet is processed.

## **flow specific**

A stream of data traveling between two endpoints across a network. Specified by source IP, destination IP, source port, destination port and IP protocol type.

## **flow table**

A table maintained on the NPM that maps individual flows to their respective processors.

## **FPGA**

Field Programmable Gate Array. A gate array where the logic network can be programmed into the device after its manufacture. An FPGA consists of an array of logic elements, either gates or lookup table RAMs, flip-flops, and programmable interconnect wiring.

## **FTP**

File Transfer Protocol.

## **gateway**

A Layer 3 devices with at least two logical interfaces, that uses a routing table to forward packets between interfaces. Note that a gateway may also act as a multi-homed host.

## **GBIC**

Gigabit Interface Converter. A transceiver that converts electric currents (digital highs and lows) to optical signals, and optical signals to digital electric currents. The GBIC is typically employed in fiber optic and Ethernet systems as an interface for high-speed networking. The data transfer rate is one gigabit per second (1 Gbps) or more.

## **GLM**

Gigabit Link Module.

## **GRUB**

GRand Unified Bootloader.

## **hash**

A cryptographic operation where an entire message is run through a mathematical operation that results in a fixed-length string that is unique.

## **HTTP**

Hypertext Transfer Protocol.

## **IDEA**

International Data Encryption Algorithm. A conventional encryption algorithm, using block cipher, operating on 64-bit blocks with a 128 bit key.

## **IDS**

Intrusion Detection System.

## **IOP**

I/O Processor.

## **IP Address**

Internet Protocol (IP). A numerical address that identifies senders and receivers of Internet data. The address accompanies data packets, identifying a particular network on the Internet and the specific device (such as a server) from which it originated.

## **IPS**

Intrusion Prevention System.

## **In-Service Upgrade**

ISU is an alternate method of upgrading XOS software while minimizing downtime. This feature has several requirements for successful completion of an ISU, including redundant CPMs, APMs, and NPMs. During ISU, the chassis is virtually split in two halves during which time only one half of the chassis will be responsible for forwarding traffic.

## **LACP**

Link Aggregation Control Protocol.

## **load balancing**

Distributing flows in real time amongst multiple APMs.

## **load table**

A table that maps flow profiles to weighted lists of virtual processors.

## **logical interface**

A channelized interface on a physical interface. A subdivision of a physical interface. Currently supported logical interface types are default and VLAN.

## **logical line**

A combination of a physical line and a sub-line (channel). A logical line is uniquely identified by a physical line ID or channel ID pair.

## **MD5**

Message Digest 5. A one-way function that takes a variable-length message and produces a fixed-length hash.

## **MAC address**

Media Access Control (MAC). A hardware address that uniquely identifies each node of a network. In IEEE 802 networks, the Data Link Control (DLC) layer of the OSI Reference Model is divided into two sub-layers: the Logical Link Control (LLC) layer and the Media Access Control (MAC) layer. The MAC layer interfaces directly with the network media. Consequently, each different type of network media requires a different MAC layer.

## **MIB**

Management Information Base.

## **MS**

Management Server.

## **multi-homed host**

A Layer 3 device with at least two logical interfaces that generate packets but does not forward the packets.

## **NMS**

Network Management System. Platform that manages one or more X-Series systems in a networked environment.

## **NPM**

Network Processor Module. The XOS module responsible for network interface access (up to 10 Gb/sec full-duplex), flow classification, distribution of flows to APMs, and load balancing of the APMs.

## **PGP**

Pretty Good Privacy. A high-security RSA public-key encryption application that enables files or messages to be exchanged with privacy and authentication.

## **physical interface**

The physical hardware connector on the NPM or CPM representing a network interface port.

## **POS**

Packet Over Sonet.

## **POST**

Power On Self-Test.

## **PPP**

Point to Point Protocol.

## **RAID**

Redundant Array of Inexpensive/Independent Drives. A data storage scheme used to allow multiple drives to work as a single drive. RAID level 0 and level 1 are supported by Crossbeam Systems in our newer modules. RAID 0 writes data to whichever drive is currently free. This method is used for greater data speed efficiency (however, all drives in the RAID are needed to fully access all the data). RAID 1 writes identical data to all the drives in the RAID grouping. This method is used for greater data integrity.

## **RDRAM**

Rambus Direct Random Access Memory.

## **RMON**

Remote Network Monitoring.

## **RPM**

Red Hat Package Manager.

## **SCP**

Secure copy.

## **SMP**

Symmetric Multi Processor.

## **SNMP**

Simple Network Management Protocol. The Internet standard protocol developed to manage and monitor nodes on an IP network.

## **SSH**

Secure Shell. A powerful authentication and encryption program replacing older and less secure tools like Telnet. SSH provides both authentication and encryption and is therefore the preferred method of network access. SSH allows a secure connection to be established between a client computer and a server host. The X-Series system provides SSH server, SSH client, and scp capability.

## **SSL**

Secure Socket Layer.

## **Stateful Inspection (dynamic packet filtering)**

A firewall architecture operating at the network layer. Unlike static packet filtering, which examines a packet based on the information in its header, stateful inspection examines not just the header information but also the contents of the packet up through the application layer in order to determine more about the packet than just information about its source and destination. A stateful inspection firewall also monitors the state of the connection and compiles the information in a state table. Because of this, filtering decisions are based not only on administrator-defined rules (as in static packet filtering) but also on context that has been established by prior packets that have passed through the firewall. As an added security measure against port scanning, stateful inspection firewalls close off ports until connection to the specific port is requested.

## **static route**

A user-defined route that causes packets to move between a source and destination along a specific path.

## **sub-line**

A multiplexed channel within a single line. Examples include: a DS0 channel within a T1/T3 serial interface, a ATM PVC, and a tagged VLAN. A sub-line is uniquely identified by a 32-bit channel ID.

## **SYSLOGD**

System Logger Daemon.

## **Telnet**

An administrator can enable Telnet as part of the boot dialogue, or by using a CLI command. Telnet comes disabled because traffic is not encrypted between the client and the X-Series system.

## **VAP**

Virtual Application Processor. An application operating environment which can be run on an APM. A VAP consists of the OS, system software, and a set of applications which run concurrently.

## **VAP group**

A virtual set of Application Processor Modules identically configured for load balancing and redundancy to process the same set of applications.

## **VLAN**

Virtual Local Area Network. A local area network with a definition that maps workstations on some other basis than geographic location (for example, by department, type of user, or primary application).

## **VND**

Virtual Network Device. A Linux kernel object representing a logical network interface. A virtual network device is directly mapped to an NPM circuit.

## **VPN**

Virtual Private Network. Consists of private lines, switching equipment and other networking equipment that are provided for the exclusive use of one customer. A VPN gives users a secure way to access resources over the Internet or other public or private networks using encryption, authentication, and tunneling.

## **VRRP**

Virtual Router Redundancy Protocol. This protocol allows several routers on a multiaccess link to utilize the same virtual IP address. One router will be elected as a master with the other routers acting as backups in case of the failure of the master router. The protocol should also support the ability to load share traffic when both routers are up.

A Virtual Router in XOS is an IP address or a set of IP addresses that can be instantiated on a circuit for a subset of the VAP groups on which the circuit is configured, and active only on one of the X-Series systems participating in multi-system High Availability configuration.

## **XML**

Extensible Markup Language. The universal format for structured documents and data on the Web as defined by a set of specifications and recommendations from the W3C.

