Network Configuration Example

Configuring Interchassis Redundancy for MX Series 3D Universal Edge Routers Using a Virtual Chassis

Published: 2014-08-11

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Network Configuration Example Configuring Interchassis Redundancy for MX Series 3D Universal Edge Routers Using a Virtual Chassis
NCE0118

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Document

This document provides an overview of the interchassis redundancy and Virtual Chassis features, and describes the benefits and components of using an MX Series Virtual Chassis to configure interchassis redundancy in your network. The document includes a detailed example that describes how to set up and configure an MX Series Virtual Chassis consisting of two Juniper Networks® MX Series 3D Universal Edge Routers, each with dual Routing Engines and Modular Port Concentrators (MPCs) installed, for Junos OS Release 11.2 and later releases.

Interchassis Redundancy and Virtual Chassis Overview

As more high-priority voice and video traffic is carried on the network, interchassis redundancy has become a baseline requirement for providing stateful redundancy on broadband subscriber management equipment such as broadband services routers, broadband network gateways, and broadband remote access servers. To provide a stateful interchassis redundancy solution for MX Series 3D Universal Edge Routers, you can configure a Virtual Chassis.

This topic provides an overview of interchassis redundancy and the Virtual Chassis, and explains the benefits of configuring a Virtual Chassis on supported MX Series routers.

Interchassis Redundancy Overview

Traditionally, redundancy in broadband edge equipment has used an intrachassis approach, which focuses on providing redundancy within a single system. However, a single-system redundancy mechanism no longer provides the degree of high availability required by service providers who must carry mission-critical voice and video traffic on their network. Consequently, service providers are requiring interchassis redundancy solutions that can span multiple systems that are colocated or geographically dispersed.

Interchassis redundancy is a high availability feature that prevents network outages and protects routers against access link failures, uplink failures, and wholesale chassis failures without visibly disrupting the attached subscribers or increasing the network management burden for service providers. Network outages can cause service providers to lose revenues and require them to register formal reports with government agencies. A robust interchassis redundancy implementation enables service providers to fulfill strict service-level agreements (SLAs) and avoid unplanned network outages to better meet the needs of their customers.

Virtual Chassis Overview

One approach to providing interchassis redundancy is the Virtual Chassis model. In general terms, a *Virtual Chassis* configuration enables a collection of member routers to function as a single virtual router, and extends the features available on a single router to the member routers in the Virtual Chassis. The interconnected member routers in a Virtual Chassis are managed as a single network element that appears to the network administrator as a single chassis with additional line card slots, and to the access network as a single system.

To provide a stateful interchassis redundancy solution for MX Series routers, you can configure a Virtual Chassis. An MX Series Virtual Chassis interconnects two MX Series routers into a logical system that you can manage as a single network element. The member routers in a Virtual Chassis are designated as the *Virtual Chassis master router* (also known as the *protocol master*) and the *Virtual Chassis backup router* (also known as the *protocol backup*). The member routers are interconnected by means of dedicated *Virtual Chassis ports* that you configure on Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces.

An MX Series Virtual Chassis is managed by the *Virtual Chassis Control Protocol* (*VCCP*), which is a dedicated control protocol based on IS-IS. VCCP runs on the Virtual Chassis port interfaces and is responsible for building the Virtual Chassis topology, electing the Virtual Chassis master router, and establishing the interchassis routing table to route traffic within the Virtual Chassis.

Supported Platforms for MX Series Virtual Chassis

You can configure a Virtual Chassis on the following MX Series routers with MPC/MIC interfaces (for configuration of Virtual Chassis ports) and dual Routing Engines:

- MX240 3D Universal Edge Router
- MX480 3D Universal Edge Router
- MX960 3D Universal Edge Router

In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled on both member routers in the Virtual Chassis.

Benefits of Configuring a Virtual Chassis

Configuring a Virtual Chassis for MX Series routers provides the following benefits:

- Simplifies network management of two routers that are either colocated or geographically dispersed across a Layer 2 point-to-point network.
- Provides resiliency against network outages and protects member routers against access link failures, uplink failures, and chassis failures without visibly disrupting attached subscribers or increasing the network management burden for service providers.
- Extends the high availability capabilities of applications such as GRES and NSR beyond a single MX Series router to both member routers in the Virtual Chassis.
- Enables service providers to fulfill strict service level agreements (SLAs) and avoid unplanned network outages to better meet their customers' needs.
- Provides the ability to scale bandwidth and service capacity as more high-priority voice and video traffic is carried on the network.

Related Documentation

- Virtual Chassis Components Overview on page 3
- Example: Configuring Interchassis Redundancy for MX Series Routers Using a Virtual Chassis on page 9

Virtual Chassis Components Overview

A Virtual Chassis configuration for MX Series 3D Universal Edge Routers interconnects two MX Series routers into a logical system that you can manage as a single network element.Figure 1 on page 3 illustrates a typical topology for a two-member MX Series Virtual Chassis.

Figure 1: Sample Topology for MX Series Virtual Chassis

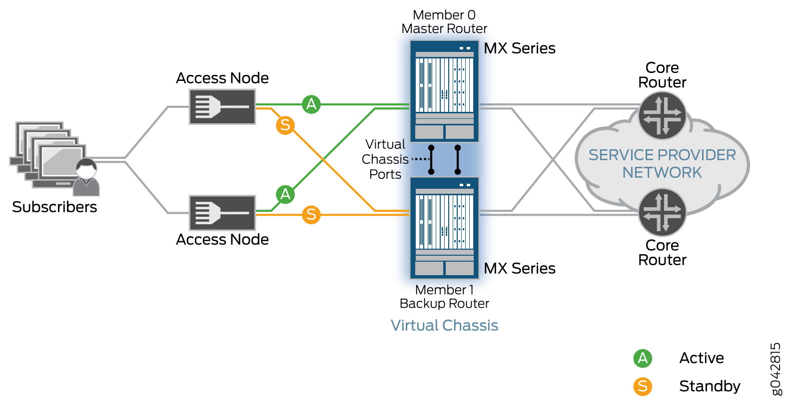

This overview describes the basic hardware and software components of the Virtual Chassis configuration illustrated in Figure 1 on page 3, and covers the following topics:

Virtual Chassis Master Router

One of the two member routers in the Virtual Chassis becomes the *master router*, also known as the *protocol master*. The Virtual Chassis master router maintains the global configuration and state information for both member routers, and runs the chassis management processes. The master Routing Engine that resides in the Virtual Chassis master router becomes the global master for the Virtual Chassis.

Specifically, the master Routing Engine that resides in the Virtual Chassis master router performs the following functions in a Virtual Chassis:

- Manages both the master and backup member routers
- Runs the chassis management processes and control protocols
- Receives and processes all incoming and exception path traffic destined for the Virtual Chassis
- Propagates the Virtual Chassis configuration (including member IDs, roles, and configuration group definitions and applications) to the members of the Virtual Chassis

The first member of the Virtual Chassis becomes the initial master router by default. After the Virtual Chassis is formed with both member routers, the Virtual Chassis Control Protocol (VCCP) software runs a mastership election algorithm to elect the master router for the Virtual Chassis configuration.

NOTE: **You cannot configure mastership election for an MX Series Virtual Chassis.**

Virtual Chassis Backup Router

The member router in the Virtual Chassis that is not designated as the master router becomes the *backup router*, also known as the *protocol backup*. The Virtual Chassis backup router takes over mastership of the Virtual Chassis if the master router is unavailable, and synchronizes routing and state information with the master router. The master Routing Engine that resides in the Virtual Chassis backup router becomes the global backup for the Virtual Chassis.

Specifically, the master Routing Engine that resides in the Virtual Chassis backup router performs the following functions in a Virtual Chassis:

- If the master router fails or is unavailable, takes over mastership of the Virtual Chassis in order to preserve routing information and maintain network connectivity without disruption
- Synchronizes routing and application state, including routing tables and subscriber state information, with the master Routing Engine that resides in the Virtual Chassis master router
- Relays chassis control information, such as line card presence and alarms, to the master router

Virtual Chassis Line-Card Role

NOTE: **The line-card role is not supported in the preprovisioned configuration for a two-member MX Series Virtual Chassis. The line-card role applies only in the context of split detection behavior.**

A member router functioning in the **line-card** role runs only a minimal set of chassis management processes required to relay chassis control information, such as line card presence and alarms, to the Virtual Chassis master router.

You cannot explicitly configure a member router with the **line-card** role. However, if the backup router fails in a two-member Virtual Chassis configuration and split detection is enabled (the default behavior), the master router takes a **line-card** role, and line cards (FPCs) that do not host Virtual Chassis ports go offline. This state effectively isolates the master router and removes it from the Virtual Chassis until connectivity is restored. As a result, routing is halted and the Virtual Chassis configuration is disabled.

Virtual Chassis Ports

Virtual Chassis ports are special Ethernet interfaces that form a point-to-point connection between the member routers in a Virtual Chassis. When you create a Virtual Chassis, you must configure the Virtual Chassis ports on Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces. After you configure a Virtual Chassis port, it is renamed **vcp-*slot/pic/port*** (for example, **vcp-2/2/0**), and the line card associated with that port comes online. For example, the sample Virtual Chassis topology shown in Figure 1 on page 3 has a total of four Virtual Chassis ports (represented by the dots), two on each of the two member routers.

After a Virtual Chassis port is configured, it is dedicated to the task of interconnecting member routers, and is no longer available for configuration as a standard network port. To restore this port to the global configuration and make it available to function as a standard network port, you must delete the Virtual Chassis port from the Virtual Chassis configuration.

NOTE: The Junos® operating system (Junos OS) software enables you to preconfigure ports that are currently unavailable for use. Although a Virtual Chassis port is unavailable for use as a standard network port, you can configure this port as a standard network port even after you configure it as a Virtual Chassis port. However, the router does not apply the configuration until you delete the Virtual Chassis port from the Virtual Chassis configuration.

You can configure a Virtual Chassis port on either a 1-Gigabit Ethernet (**ge**) interface or a 10-Gigabit Ethernet (**xe**) interface. However, you cannot configure a combination of 1-Gigabit Ethernet Virtual Chassis ports and 10-Gigabit Ethernet Virtual Chassis ports in the same Virtual Chassis. We recommend that you configure Virtual Chassis ports only on 10-Gigabit Ethernet interfaces. In addition, to minimize network disruption in the event of a router or link failure, configure redundant Virtual Chassis ports that reside on different line cards in each member router.

Virtual Chassis port interfaces carry both Virtual Chassis Control Protocol (VCCP) packets and internal control and data traffic. Because the internal control traffic is neither encrypted nor authenticated, make sure that the Virtual Chassis port interfaces are properly secured to prevent malicious third-party attacks on the data.

Virtual Chassis ports use a default class-of-service (CoS) configuration that applies equally to all Virtual Chassis port interfaces configured in a Virtual Chassis. Optionally, you can create a customized CoS traffic-control profile and apply it to all Virtual Chassis port interfaces. For example, you might want to create a nondefault traffic-control profile that allocates more than the default 5 percent of the Virtual Chassis port bandwidth to control traffic, or that assigns different priorities and excess rates to different forwarding classes.

Virtual Chassis Port Trunks

If two or more Virtual Chassis ports of the same type and speed are configured between the same two member routers in an MX Series Virtual Chassis, the VCCP software bundles these Virtual Chassis port interfaces into a trunk, reduces the routing cost accordingly, and performs traffic load balancing across all of the Virtual Chassis port interfaces (also referred to as Virtual Chassis port links) in the trunk.

A Virtual Chassis port trunk must include only Virtual Chassis ports of the same type and speed. For example, a Virtual Chassis port trunk can include either all 10-Gigabit Ethernet (**xe** media type) Virtual Chassis ports or all 1-Gigabit Ethernet (**ge** media type) Virtual Chassis ports. An MX Series Virtual Chassis does *not* support a combination of 1-Gigabit Ethernet Virtual Chassis ports and 10-Gigabit Ethernet Virtual Chassis ports in the same Virtual Chassis port trunk.

The router uses the following formula to determine the cost metric of a Virtual Chassis port link in a Virtual Chassis port trunk:

Cost = (300 * 1,000,000,000) / *port-speed*

where *port-speed* is the aggegate speed, in bits per second, of the Virtual Chassis port.

For example, a 10-Gigabit Ethernet Virtual Chassis port link has a cost metric of 30 (300 * 1,000,000,000 / 10,000,000,000). A 1-Gigabit Ethernet Virtual Chassis port link has a cost metric of 300 (300 * 1,000,000,000 / 1,000,000,000). Virtual Chassis port links with a lower cost metric are preferred over those with a higher cost metric.

An MX Series Virtual Chassis supports up to 16 Virtual Chassis ports per trunk.

Slot Numbering in the Virtual Chassis

When the Virtual Chassis forms, the slots for line cards (FPCs) that do not host Virtual Chassis ports are renumbered to reflect the slot numbering and offsets used in the Virtual Chassis instead of the physical slot numbers where the line card is actually installed. In a two-member MX Series Virtual Chassis, member 0 in the Virtual Chassis uses FPC slot numbers 0 through 11 with no offset, and member 1 uses FPC slot numbers 12 through 23, with an offset of 12.

For example, a 10-Gigabit Ethernet interface that appears as **xe-14/2/2** (FPC slot 14, PIC slot 2, port 2) in the **show interfaces** command output is actually physical interface **xe-2/2/2** (FPC slot 2, PIC slot 2, port 2) on member 1 after deducting the FPC slot numbering offset of 12 for member 1.

The slot numbering for Virtual Chassis ports uses the physical slot number where the Virtual Chassis port is configured. For example, **vcp-3/2/0** is configured on physical FPC slot 3, PIC slot 2, port 0.

Virtual Chassis Control Protocol

An MX Series Virtual Chassis is managed by the Virtual Chassis Control Protocol (VCCP), which is a dedicated control protocol based on IS-IS. VCCP runs on the Virtual Chassis port interfaces and performs the following functions in the Virtual Chassis:

- Discovers and builds the Virtual Chassis topology
- Runs the mastership election algorithm to determine the Virtual Chassis master router
- Establishes the interchassis routing table to route traffic within the Virtual Chassis

Like IS-IS, VCCP exchanges link-state PDUs for each member router to construct a shortest path first (SPF) topology and to determine each member router's role (master or backup) in the Virtual Chassis. Because VCCP supports only point-to-point connections, no more than two member routers can be connected on any given Virtual Chassis port interface.

Member IDs, Roles, and Serial Numbers

To configure an MX Series Virtual Chassis, you must create a preprovisioned configuration that provides the following required information for each member router:

- Member ID—A numeric value (**0** or **1**) that identifies the member router in a Virtual Chassis configuration.
- Role—The role to be performed by each member router in the Virtual Chassis. In a two-member MX Series Virtual Chassis, you must assign both member routers the **routing-engine** role, which enables either router to function as the master router or backup router of the Virtual Chassis.
- Serial number—The chassis serial number of each member router in the Virtual Chassis. To obtain the router's serial number, find the label affixed to the side of the MX Series chassis, or issue the **show chassis hardware** command on the router to display the serial number in the command output.

The preprovisioned configuration permanently associates the member ID and role with the member router's chassis serial number. When a new member router joins the Virtual Chassis, the VCCP software compares the router's serial number against the values specified in the preprovisioned configuration. If the serial number of a joining router does not match any of the configured serial numbers, the VCCP software prevents that router from becoming a member of the Virtual Chassis.

Related Documentation

- Interchassis Redundancy and Virtual Chassis Overview on page 1
- Guidelines for Configuring Virtual Chassis Ports on page 8
- Example: Configuring Interchassis Redundancy for MX Series Routers Using a Virtual Chassis on page 9

Guidelines for Configuring Virtual Chassis Ports

To interconnect the member routers in a Virtual Chassis for MX Series 3D Universal Edge Routers, you must configure Virtual Chassis ports on Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces. After you configure a Virtual Chassis port, that port is dedicated to the task of interconnecting member routers, and is no longer available for configuration as a standard network port.

NOTE: **The Junos OS software enables you to preconfigure ports that are currently unavailable for use. Although a Virtual Chassis port is unavailable for use as a standard network port, you can configure this port as a standard network port even after you configure it as a Virtual Chassis port. However, the router does not apply the configuration until you delete the Virtual Chassis port from the Virtual Chassis configuration.**

Consider the following guidelines when you configure Virtual Chassis ports in an MX Series Virtual Chassis:

- An MX Series Virtual Chassis supports up to 16 Virtual Chassis ports per trunk.

  If two or more Virtual Chassis ports of the same type and speed (that is, either all 10-Gigabit Ethernet Virtual Chassis ports or all 1-Gigabit Ethernet Virtual Chassis ports) are configured between the same two member routers in an MX Series Virtual Chassis, the Virtual Chassis Control Protocol (VCCP) bundles these Virtual Chassis port interfaces into a trunk, reduces the routing cost accordingly, and performs traffic load balancing across all of the Virtual Chassis port interfaces in the trunk.

- An MX Series Virtual Chassis does *not* support a combination of 1-Gigabit Ethernet (**ge** media type) Virtual Chassis ports and 10-Gigabit Ethernet (**xe** media type) Virtual Chassis ports within the same Virtual Chassis.

  You must configure either all 10-Gigabit Virtual Chassis ports or all 1-Gigabit Virtual Chassis ports in the same Virtual Chassis. We recommend that you configure Virtual Chassis ports on 10-Gigabit Ethernet (**xe**) interfaces.

  This restriction has no effect on access ports or uplink ports in an MX Series Virtual Chassis configuration.

- Configure redundant Virtual Chassis ports that reside on different line cards in each member router.

  For a two-member MX Series Virtual Chassis, we recommend that you configure a minimum of two 10-Gigabit Ethernet Virtual Chassis ports on different line cards in each member router, for a total minimum of four 10-Gigabit Ethernet Virtual Chassis ports in the Virtual Chassis. In addition, make sure that the Virtual Chassis port bandwidth is equivalent to no less than 50 percent of the aggregate bandwidth required for user data traffic. The following examples illustrate these recommendations:

  - If the bandwidth in your network is equivalent to two 10-Gigabit Ethernet interfaces (20 Gbps) on the access-facing side of the Virtual Chassis and two 10-Gigabit

Ethernet interfaces (20 Gbps) on the core-facing side of the Virtual Chassis, we recommend that you configure two 10-Gigabit Ethernet Virtual Chassis ports, which is the recommended minimum in a Virtual Chassis for redundancy purposes.

  - If the aggregate bandwidth in your network is equivalent to ten 10-Gigabit Ethernet interfaces (100 Gbps), we recommend that you configure a minimum of five 10-Gigabit Ethernet Virtual Chassis ports, which is 50 percent of the aggregate bandwidth.

- A user data packet traversing the Virtual Chassis port interfaces between member routers is discarded at the Virtual Chassis egress port if the maximum transmission unit (MTU) size of the packet exceeds 9150 bytes.

  The maximum MTU size of a 1-Gigabit Ethernet interface or 10-Gigabit Ethernet interface on a single MX Series router is 9192 bytes. In an MX Series Virtual Chassis configuration, user data packets that traverse Gigabit Ethernet or 10-Gigabit Ethernet Virtual Chassis port interfaces have 42 extra bytes of Virtual Chassis-specific header data, which reduces their maximum MTU (payload) size to 9150 bytes. The user data packet is transmitted in its entirety across the Virtual Chassis port interface. However, because packet fragmentation and reassembly are not supported on Virtual Chassis port interfaces, user data packets that exceed 9150 bytes are discarded at the Virtual Chassis egress port.

Related Documentation

- Virtual Chassis Components Overview on page 3
- Example: Configuring Interchassis Redundancy for MX Series Routers Using a Virtual Chassis on page 9

Example: Configuring Interchassis Redundancy for MX Series Routers Using a Virtual Chassis

To provide interchassis redundancy for MX Series 3D Universal Edge Routers, you can configure a Virtual Chassis. A *Virtual Chassis* configuration interconnects two MX Series routers into a logical system that you can manage as a single network element. The member routers in a Virtual Chassis are interconnected by means of Virtual Chassis ports that you configure on Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces (network ports) on each MX Series router.

This example describes how to set up and configure a Virtual Chassis consisting of two MX Series routers:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.2 and later releases
- One MX240 3D Universal Edge Router
- One MX480 3D Universal Edge Router

NOTE: This configuration example has been tested using Junos OS Release 11.2 and is assumed to work on all later releases.

See Table 1 on page 11 for information about the hardware installed in each MX Series router.

BEST PRACTICE: We recommend that you use the **commit synchronize** command throughout this procedure to save any configuration changes to the Virtual Chassis.

For an MX Series Virtual Chassis, the **force** option is the default and only behavior when you issue the **commit synchronize** command. Issuing the **commit synchronize** command for an MX Series Virtual Chassis configuration has the same effect as issuing the **commit synchronize force** command.

Overview and Topology

To configure the Virtual Chassis shown in this example, you must create a preprovisioned configuration at the **[edit virtual-chassis]** hierarchy level on the router to be designated as the protocol master of the Virtual Chassis. The preprovisioned configuration includes the serial number, member ID, and role for each member router (also known as member chassis) in the Virtual Chassis. When a new member router joins the Virtual Chassis, the software compares its serial number against the values specified in the preprovisioned configuration. If the serial number of a joining router does not match any of the configured serial numbers, the software prevents that router from becoming a member of the Virtual Chassis.

After you commit the preprovisioned configuration on the master router, you must assign the preprovisioned member IDs by issuing the **request virtual-chassis member-id set** operational mode command on each router, which causes the router to reboot. When the reboot is complete, you create one or more Virtual Chassis ports by issuing the **request virtual-chassis vc-port set** operational mode command on each router. The Virtual Chassis forms when the line cards in both member routers are back online.

This example configures a Virtual Chassis that interconnects two MX Series routers, and uses the basic topology shown in Figure 2 on page 11. For redundancy, two Virtual Chassis ports are configured on each member router.

Figure 2: Sample Topology for a Virtual Chassis with Two MX Series Routers

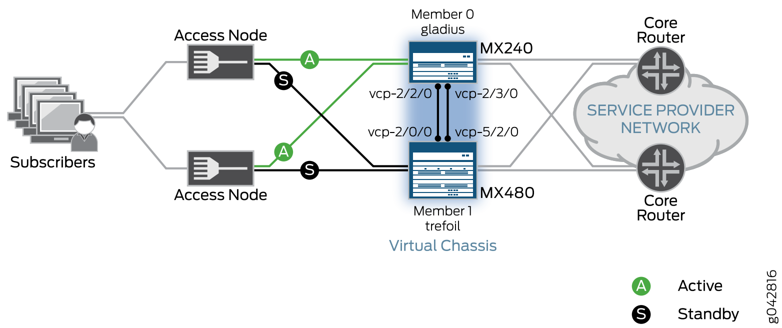

Table 1 on page 11 shows the hardware and software configuration settings for each MX Series router in the Virtual Chassis. You use some of these settings in the preprovisioned configuration and when you assign the member IDs and create the Virtual Chassis ports.

NOTE: MX Series Enhanced Queuing IP Services DPCs (DPCE-R-Q model numbers) and MX Series Enhanced Queuing Ethernet Services DPCs (DPCE-X-Q model numbers) do not interoperate with features of the MX Series Virtual Chassis. If any MX Series Enhanced Queuing DPCs are installed in a router to be configured as a member of a Virtual Chassis, you must ensure that these DPCs are offline before you configure the Virtual Chassis.

Table 1: Components of the Sample MX Series Virtual Chassis

Router Name	Hardware	Serial Number	Member ID	Role	Virtual Chassis Ports	Network Port Slot Numbering
gladius	MX240 router with: • 60-Gigabit Ethernet Enhanced Queuing MPC • 20-port Gigabit Ethernet MIC with SFP • 4-port 10-Gigabit Ethernet MIC with XFP • Master RE-S-2000 Routing Engine in slot 0 (represented in example as **member0-re0**) • Backup RE-S-2000 Routing Engine in slot 1 (represented in example as **member0-re1**)	JN10C7135AFC	0	routing-engine (master)	vcp-2/2/0 vcp-2/3/0	FPC 0 – 11

Table 1: Components of the Sample MX Series Virtual Chassis *(continued)*

Router Name	Hardware	Serial Number	Member ID	Role	Virtual Chassis Ports	Network Port Slot Numbering
trefoil	MX480 router with: • Two 30-Gigabit Ethernet Queuing MPCs • Two 20-port Gigabit Ethernet MICs with SFP • Two 2-port 10-Gigabit Ethernet MICs with XFP • Master RE-S-2000 Routing Engine in slot 0 (represented in example as **member1-re0**) • Backup RE-S-2000 Routing Engine in slot 1 (represented in example as **member1-re1**)	JN115D117AFB	1	routing-engine (backup)	vcp-2/0/0 vcp-5/2/0	FPC 12 – 23 (offset = 12)

Configuration

To configure a Virtual Chassis consisting of two MX Series routers, perform these tasks:

Preparing to Configure the Virtual Chassis

Step-by-Step Procedure

To prepare for configuring an MX Series Virtual Chassis:

1. Make a list of the serial numbers of both routers that you want to configure as part of the Virtual Chassis.

   The chassis serial number is located on a label affixed to the side of the MX Series chassis. Alternatively, you can obtain the chassis serial number by issuing the **show chassis hardware** command, which is especially useful if you are accessing the router from a remote location. For example:

```
user@gladius> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN10C7135AFC      MX240
.
.
.
Fan Tray 0       REV 01   710-021113   JT0119            MX240 Fan Tray
```

2. Note the desired role (**routing-engine**) for each router in the Virtual Chassis.

   In a two-router Virtual Chassis configuration, you must designate each router with the **routing-engine** role, which enables either router to function as the master or backup of the Virtual Chassis.

   - The *master router* maintains the global configuration and state information for all members of the Virtual Chassis, and runs the chassis management processes.
   - The *backup router* synchronizes with the master router and relays chassis control information (such as linecard presence and alarms) to the master router. If the master router is unavailable, the backup router takes mastership of the Virtual Chassis to preserve routing information and maintain network connectivity without disruption.

3. Note the member ID (0 or 1) to be assigned to each router in the Virtual Chassis.

   In this example, the master router is assigned member ID 0, and the backup router is assigned member ID 1.

4. Ensure that both MX Series routers in the Virtual Chassis have dual Routing Engines installed, and that all four Routing Engines in the Virtual Chassis are the same model.

   For example, you cannot configure a Virtual Chassis if one member router has RE-S-2000 Routing Engines installed and the other member router has RE-S-1800 Routing Engines installed.

5. Ensure that the necessary Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces on which to configure the Virtual Chassis ports are installed and operational in each member router.

NOTE: An MX Series Virtual Chassis does not support a combination of 1-Gigabit Ethernet (ge media type) Virtual Chassis ports and 10-Gigabit Ethernet (xe media type) Virtual Chassis ports within the same Virtual Chassis. You must configure either all 10-Gigabit Ethernet Virtual Chassis ports or all 1-Gigabit Ethernet Virtual Chassis ports in the same Virtual Chassis. We recommend that you configure Virtual Chassis ports on 10-Gigabit Ethernet interfaces. This restriction has no effect on access ports or uplink ports in an MX Series Virtual Chassis configuration.

6. If MX Series Enhanced Queuing IP Services DPCs (DPCE-R-Q model numbers) or MX Series Enhanced Queuing Ethernet Services DPCs (DPCE-X-Q model numbers) are installed in a router to be configured as a member of the Virtual Chassis, make sure that these DPCs are offline before you configure the Virtual Chassis. Otherwise, the MX Series Virtual Chassis configuration will not function.

   NOTE: MX Series Enhanced Queuing IP Services DPCs (DPCE-R-Q model numbers) and MX Series Enhanced Queuing Ethernet Services DPCs (DPCE-X-Q model numbers) do not interoperate with features of the MX Series Virtual Chassis.

7. Determine the desired location of the dedicated Virtual Chassis ports on both member routers, and use the Virtual Chassis ports to physically interconnect the member routers in a point-to-point topology.
8. Ensure that both MX Series routers to be configured as members of the Virtual Chassis are running the same Junos OS release, and have basic network connectivity.
9. Install the MX Virtual Chassis Redundancy Feature Pack license on each router to be configured as part of the Virtual Chassis.
10. Install the necessary Junos OS feature licenses on each router to be configured as part of the Virtual Chassis.

Creating and Applying Configuration Groups for the Virtual Chassis

Step-by-Step Procedure

For a Virtual Chassis configuration consisting of two MX Series routers, each of which supports dual Routing Engines, you must create and apply the following configuration groups on the router to be designated as the master of the Virtual Chassis instead of using the standard re0 and re1 configuration groups:

- **member0-re0**
- **member0-re1**
- **member1-re0**
- **member1-re1**

NOTE: The **member*n*-re*n*** naming format for configuration groups is reserved for exclusive use by member routers in an MX Series Virtual Chassis configuration.

To create and apply configuration group information for the Virtual Chassis:

1. Log in to the console on member 0 (**gladius**).
2. In the console window on member 0, create and apply the **member0-re0** configuration group.

```
[edit]
user@gladius# copy groups re0 to member0-re0
user@gladius# set apply-groups member0-re0
```

3. Delete the standard **re0** configuration group from the global configuration on member 0.

```
[edit]
user@gladius# delete apply-groups re0
user@gladius# delete groups re0
```

4. Create and apply the **member0-re1** configuration group on member 0.

```
[edit]
user@gladius# copy groups re1 to member0-re1
user@gladius# set apply-groups member0-re1
```

5. Delete the standard **re1** configuration group from the global configuration on member 0.

```
[edit]
user@gladius# delete apply-groups re1
user@gladius# delete groups re1
```

6. Create and apply the **member1-re0** configuration information on member 0.

```
[edit]
user@gladius# set groups member1-re0 system host-name trefoil
user@gladius# set groups member1-re0 system backup-router 10.9.0.1
```

```
user@gladius# set groups member1-re0 system backup-router destination
  172.16.0.0/12
user@gladius# set groups member1-re0 system backup-router destination
  10.9.0.0/16
...
user@gladius# set groups member1-re0 interfaces fxp0 unit 0 family inet address
  10.9.3.97/21
user@gladius# set apply-groups member1-re0
```

The examples in Steps 6 and 7 set the IP address for the **fxp0** management interface and add an IP route for it in the event that routing becomes inactive.

7. Create and apply the **member1-re1** configuration information on member 0.

```
[edit]
user@gladius# set groups member1-re1 system host-name trefoil
user@gladius# set groups member1-re1 system backup-router 10.9.0.1
user@gladius# set groups member1-re1 system backup-router destination
  172.16.0.0/12
user@gladius# set groups member1-re1 system backup-router destination 10.9.0.0/16
...
user@gladius# set groups member1-re1 interfaces fxp0 unit 0 family inet address
  10.9.3.98/21
user@gladius# set apply-groups member1-re1
```

8. Commit the configuration on member 0.

Results Display the results of the configuration.

```
[edit]
user@gladius# show groups ?
  Possible completions:
   <[Enter]>        Execute this command
   <group_name>       Group name
   global          Group name
   member0-re0       Group name
   member0-re1       Group name
   member1-re0       Group name
   member1-re1       Group name
   |             Pipe through a command

[edit]
user@gladius# show apply-groups
  apply-groups [ global member0-re0 member0-re1 member1-re0 member1-re1 ];
```

Configuring Preprovisioned Member Information for the Virtual Chassis

Step-by-Step Procedure

To configure the preprovisioned member information on member 0 (**gladius**):

1. Log in to the console on member 0.
2. Specify that you want to create a preprovisioned Virtual Chassis configuration.

```
[edit virtual-chassis]
user@gladius# set preprovisioned
```

3. Configure the member ID (**0** or **1**), role (**routing-engine**), and chassis serial number for each member router in the Virtual Chassis.


```
[edit virtual-chassis]
user@gladius# set member 0 role routing-engine serial-number JN10C7135AFC
user@gladius# set member 1 role routing-engine serial-number JN115D117AFB
```

4. Disable detection of a split in the Virtual Chassis configuration.
   By default, split detection in an MX Series Virtual Chassis is enabled.

```
[edit virtual-chassis]
user@gladius# set no-split-detection
```

BEST PRACTICE: We recommend that you disable split detection for a two-member MX Series Virtual Chassis configuration if you think the backup router is more likely to fail than the Virtual Chassis port links to the backup router. Configuring redundant Virtual Chassis ports on different line cards in each member router reduces the likelihood that all Virtual Chassis port links to the backup router will fail.

5. (Optional) Enable tracing of Virtual Chassis operations.

```
[edit virtual-chassis]
user@gladius# set traceoptions file vccp
user@gladius# set traceoptions file size 100m
user@gladius# set traceoptions flag all
```

6. Commit the configuration.

Results Display the results of the configuration.

```
[edit virtual-chassis]
user@gladius# show
preprovisioned;
no-split-detection;
traceoptions {
  file vccp size 100m;
  flag all;
}
member 0 {
  role routing-engine;
  serial-number JN10C7135AFC;
}
member 1 {
  role routing-engine;
  serial-number JN115D117AFB;
}
```

Configuring Enhanced IP Network Services

Step-by-Step Procedure

For an MX Series Virtual Chassis to function properly, you must configure enhanced IP network services on both member routers (member 0 and member 1) when you first set up the Virtual Chassis. Enhanced IP network services define how the chassis recognizes and uses certain modules. When you set each member router's network services to **enhanced-ip**, only MPC/MIC modules and MS-DPC modules are powered on in the chassis. Nonservice DPCs do not work with enhanced IP network services.

To configure enhanced IP network services for a Virtual Chassis:

1. Configure enhanced IP network services on member 0 (**gladius**).
   a. Log in to the console on member 0.
   b. Access the chassis hierarchy.

      ```
      [edit]
      user@gladius# edit chassis
      ```

   c. Configure enhanced IP network services for member 0.

      ```
      [edit chassis]
      user@gladius# set network-services enhanced-ip
      ```

   d. Commit the configuration on member 0.

      NOTE: Immediately after you commit the configuration, the software prompts you to reboot the router. You can proceed without rebooting the router at this point because a reboot occurs when you configure the member IDs to enable Virtual Chassis mode, later in this procedure.

2. Configure enhanced IP network services on member 1 (**trefoil**).
   a. Log in to the console on member 1.
   b. Access the chassis hierarchy.

      ```
      [edit]
      user@trefoil# edit chassis
      ```

   c. Configure enhanced IP network services for member 1.

      ```
      [edit chassis]
      user@trefoil# set network-services enhanced-ip
      ```

   d. Commit the configuration on member 1.

      NOTE: Immediately after you commit the configuration, the software prompts you to reboot the router. You can proceed without rebooting the router at this point because a reboot occurs when you configure the member IDs to enable Virtual Chassis mode, later in this procedure.

Enabling Graceful Routing Engine Switchover and Nonstop Active Routing

Step-by-Step Procedure

Before you configure member IDs and Virtual Chassis ports, you must enable graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) on both member routers in the Virtual Chassis.

To enable GRES and NSR:

1. Enable GRES and NSR on member 0 (**gladius**):

   a. Log in to the console on member 0.

   b. Enable graceful switchover.

      [edit chassis redundancy]
      user@gladius# **set graceful-switchover**

   c. Enable nonstop active routing.

      [edit routing-options]
      user@gladius# **set nonstop-routing**

   d. Configure the **commit** command to automatically result in a **commit synchronize** action between the dual Routing Engines in member 0.

      [edit system]
      user@gladius# **set commit synchronize**

   e. Commit the configuration.

2. Enable GRES and NSR on member 1 (**trefoil**):

   a. Log in to the console on member 1.

   b. Enable graceful switchover.

      [edit chassis redundancy]
      user@trefoil# **set graceful-switchover**

   c. Enable nonstop active routing.

      [edit routing-options]
      user@trefoil# **set nonstop-routing**

   d. Configure the **commit** command to automatically result in a **commit synchronize** action between the dual Routing Engines in member 1.

      [edit system]
      user@trefoil# **set commit synchronize**

   e. Commit the configuration.

NOTE: When you configure nonstop active routing, you must include the **commit synchronize** statement at the **[edit system]** hierarchy level. Otherwise, the commit operation fails.

For an MX Series Virtual Chassis, the **force** option is the default and only behavior when you use the **commit synchronize** statement. Including the **commit synchronize** statement for an MX Series Virtual Chassis configuration has the same effect as including the **commit synchronize force** statement.

Configuring Member IDs and Rebooting the Routers to Enable Virtual Chassis Mode

Step-by-Step Procedure

To configure (set) the preprovisioned member ID for each MX Series router in the Virtual Chassis, use the **request virtual-chassis member-id set** command. Assigning the member ID causes the router to reboot in preparation for forming the Virtual Chassis.

NOTE: If you issue the **request virtual-chassis member-id set** command without first installing an MX Virtual Chassis Redundancy Feature Pack license on both member routers, the software displays a warning message that you are operating without a valid Virtual Chassis software license.

To configure the member ID and reboot each router to enable Virtual Chassis mode:

1. Log in to the console on member 0 (**gladius**).
2. Set the member ID on member 0.

   user@gladius> **request virtual-chassis member-id set member 0**

   ```
   This command will enable virtual-chassis mode and reboot the system.
    Continue? [yes,no] yes
   ```

   The router reboots in preparation for membership in the Virtual Chassis. After the reboot, all MPCs remain powered off until the Virtual Chassis port connection is configured.
3. Log in to the console on member 1 (**trefoil**).
4. Set the member ID on member 1.

   user@trefoil> **request virtual-chassis member-id set member 1**

   ```
   This command will enable virtual-chassis mode and reboot the system.
    Continue? [yes,no] yes
   ```

   After the reboot, all MPCs remain powered off until the Virtual Chassis port connection is configured.

Results Display the results of the configuration on each router. At this point in the procedure, all line cards are offline, and the routers are each designated with the **Master** role because they are not yet interconnected as a fully formed Virtual Chassis. In addition, member 1 (**trefoil**) remains in **Amnesiac** state (has no defined configuration) until the Virtual Chassis forms and you commit the configuration.

For member 0 (**gladius**):

```
{master:member0-re0}

user@gladius> show virtual-chassis status
Preprovisioned Virtual Chassis
Virtual Chassis ID: 4f2b.1aa0.de08
                                               Mastership          Neighbor List

Member ID       Status   Serial No    Model    priority    Role      ID  Interface

0 (FPC   0- 11) Prsnt    JN10C7135AFC mx240         129  Master*
```

For member 1 (**trefoil**):

```
Amnesiac (ttyd0)

login: user
Password:
...

{master:member1-re0}

user@trefoil> show virtual-chassis status
Virtual Chassis ID: eabf.4e50.91e6
Virtual Chassis Mode: Disabled
                                               Mastership          Neighbor List

Member ID       Status   Serial No    Model    priority    Role      ID  Interface

1 (FPC  12- 23) Prsnt    JN115D117AFB mx480         128  Master*
```

Configuring Virtual Chassis Ports to Interconnect Member Routers

Step-by-Step Procedure

To interconnect the member routers in an MX Series Virtual Chassis, use the **request virtual-chassis vc-port set** command to configure (set) Virtual Chassis ports on Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces.

NOTE: **If you issue the request virtual-chassis vc-port set command without first installing an MX Virtual Chassis Redundancy Feature Pack license on both member routers, the software displays a warning message that you are operating without a valid Virtual Chassis software license.**

To configure Virtual Chassis ports on MPC/MIC interfaces to connect the member routers in the Virtual Chassis:

1. Configure the Virtual Chassis ports on member 0 (**gladius**).

   a. Log in to the console on member 0.

   b. Configure the first Virtual Chassis port that connects to member 1 (**trefoil**).

   ```
   {master:member0-re0}
   user@gladius> request virtual-chassis vc-port set fpc-slot 2 pic-slot 2 port 0
   vc-port successfully set
   ```

   After the Virtual Chassis port is created, it is renamed **vcp-*slot/pic/port*** (for example, **vcp-2/2/0**), and the line card associated with that port comes online. The line cards in the other member router remain offline until the Virtual Chassis forms. Each Virtual Chassis port is dedicated to the task of interconnecting member routers in a Virtual Chassis, and is no longer available for configuration as a standard network port.

   c. When **vcp-2/2/0** is up, configure the second Virtual Chassis port that connects to member 1.

   ```
   {master:member0-re0}
   user@gladius> request virtual-chassis vc-port set fpc-slot 2 pic-slot 3 port 0
   vc-port successfully set
   ```

2. Configure the Virtual Chassis ports on member 1 (**trefoil**).

   a. Log in to the console on member 1.

   b. Configure the first Virtual Chassis port that connects to member 0 (**gladius**).

   ```
   {master:member1-re0}
   user@trefoil> request virtual-chassis vc-port set fpc-slot 2 pic-slot 0 port 0
   vc-port successfully set
   ```

   c. When **vcp-2/0/0** is up, configure the second Virtual Chassis port that connects to member 0.

   ```
   {master:member1-re0}
   user@trefoil> request virtual-chassis vc-port set fpc-slot 5 pic-slot 2 port 0
   ```

```
vc-port successfully set
```

When all of the line cards in all of the member routers are online and the Virtual Chassis forms, you can issue Virtual Chassis commands from the terminal window of the master router (**gladius**).

3. Verify that the Virtual Chassis is properly configured and operational.

```
{master:member0-re0}
user@gladius> show virtual-chassis status
{master:member0-re0}
user@gladius> show virtual-chassis vc-port all-members
```

See the Verification section for information about interpreting the output of these commands.

4. Commit the configuration on the master router.

The commit step is required to ensure that the configuration groups and Virtual Chassis configuration are propagated to both members of the Virtual Chassis.

Verification

Confirm that the configuration is working properly.

- Verifying the Member IDs and Roles of the Virtual Chassis Members on page 23
- Verifying the Enhanced IP Network Services Configuration on page 24
- Verifying the Operation of the Virtual Chassis Ports on page 24
- Verifying Neighbor Reachability on page 25

Verifying the Member IDs and Roles of the Virtual Chassis Members

Purpose Verify that the member IDs and roles of the routers belonging to the Virtual Chassis are properly configured.

Action Display the status of the members of the Virtual Chassis configuration.

```
{master:member0-re0}
user@gladius> show virtual-chassis status
Preprovisioned Virtual Chassis
Virtual Chassis ID: a5b6.be0c.9525
                                              Mastership           Neighbor List
Member ID        Status    Serial No    Model    priority    Role      ID  Interface

0 (FPC   0- 11) Prsnt     JN10C7135AFC mx240          129   Master*    1  vcp-2/2/0
                                                                       1  vcp-2/3/0

1 (FPC  12- 23) Prsnt     JN115D117AFB mx480          129   Backup     0  vcp-2/0/0
                                                                       0  vcp-5/2/0
```

Meaning The value **Prsnt** in the **Status** column of the output confirms that the member routers specified in the preprovisioned configuration are currently connected to the Virtual Chassis. The display shows that member 0 (**gladius**) and member 1 (**trefoil**), which were both configured with the **routing-engine** role, are functioning as the master router and backup router of the Virtual Chassis, respectively. The **Neighbor List** displays the interconnections between the member routers by means of the Virtual Chassis ports. For example, member 0 is connected to member 1 through **vcp-2/2/0** and **vcp-2/3/0**. The asterisk (*) following **Master** denotes the router on which the command was issued. The **Mastership priority** value is assigned by the software and is not configurable.

Verifying the Enhanced IP Network Services Configuration

Purpose Verify that enhanced IP network services have been properly configured for the Virtual Chassis.

Action Display the setting of the network services configuration for the master Routing Engine in the Virtual Chassis master router (member0-re0), and for the master Routing Engine in the Virtual Chassis backup router (member1-re0).

```
{master:member0-re0}
user@gladius> show chassis network services
Network Services Mode: Enhanced-IP

{backup:member1-re0}
user@trefoil> show chassis network services
Network Services Mode: Enhanced-IP
```

Meaning The output of the **show chassis network services** command confirms that enhanced IP network services are properly configured on both member routers in the Virtual Chassis.

Verifying the Operation of the Virtual Chassis Ports

Purpose Verify that the Virtual Chassis ports are properly configured and operational.

Action Display the status of the Virtual Chassis ports for both members of the Virtual Chassis.

```
{master:member0-re0}

user@gladius> show virtual-chassis vc-port all-members
member0:
--------------------------------------------------------------------------
Interface        Type              Trunk  Status       Speed        Neighbor
or                                  ID                 (mbps)       ID  Interface
Slot/PIC/Port
2/2/0            Configured          3    Up           10000        1   vcp-2/0/0
2/3/0            Configured          3    Up           10000        1   vcp-5/2/0


member1:
--------------------------------------------------------------------------
Interface        Type              Trunk  Status       Speed        Neighbor
or                                  ID                 (mbps)       ID  Interface
Slot/PIC/Port
2/0/0            Configured          3    Up           10000        0   vcp-2/2/0
5/2/0            Configured          3    Up           10000        0   vcp-2/3/0
```

Meaning The output confirms that the Virtual Chassis ports you configured are operational. For each member router, the **Interface or Slot/PIC/Port** column shows the location of the Virtual Chassis ports configured on that router. For example, the Virtual Chassis ports on **member0-re0** (**gladius**) are **vcp-2/2/0** and **vcp-2/3/0**. In the **Trunk ID** column, the value **3** indicates that a trunk has formed; if a trunk is not present, this field displays the value **-1** . In the **Status** column, the value **Up** confirms that the interfaces associated with the Virtual Chassis ports are operational. The **Speed** column displays the speed of the Virtual Chassis port interface. The **Neighbor ID/Interface** column displays the member IDs and Virtual Chassis port interfaces that connect to this router. For example, the connections to member 0 (**gladius**) are through **vcp-2/0/0** and **vcp-5/2/0** on member 1 (**trefoil**).

Verifying Neighbor Reachability

Purpose Verify that each member router in the Virtual Chassis can reach the neighbor routers to which it is connected.

Action Display the neighbor reachability information for both member routers in the Virtual Chassis.

```
{master:member0-re0}
user@gladius> show virtual-chassis active-topology all-members
member0:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  1                    1(vcp-2/2/0.32768)


member1:
--------------------------------------------------------------------------
  Destination ID        Next-hop

  0                    0(vcp-2/0/0.32768)
```

Meaning The output confirms that each member router in the Virtual Chassis has a path to reach the neighbors to which it is connected. For each member router, the **Destination ID** specifies the member ID of the destination (neighbor) router. The **Next-hop** column displays the member ID and Virtual Chassis port interface of the next hop to which packets for the destination ID are forwarded. For example, the next hop from member 0 (**gladius**) to member 1 (**trefoil**) is through Virtual Chassis port interface **vcp-2/2/0.32768**.

Related Documentation

- Interchassis Redundancy and Virtual Chassis Overview on page 1
- Virtual Chassis Components Overview on page 3
- Guidelines for Configuring Virtual Chassis Ports on page 8

