

VIRTUAL CHASSIS TECHNOLOGY BEST PRACTICES

Although Juniper Networks has attempted to provide accurate information in this guide, Juniper Networks does not warrant or guarantee the accuracy of the information provided herein. Third party product descriptions and related technical details provided in this document are for information purposes only and such products are not supported by Juniper Networks. All information provided in this guide is provided "as is", with all faults, and without warranty of any kind, either expressed or implied or statutory. Juniper Networks and its suppliers hereby disclaim all warranties related to this guide and the information contained herein, whether expressed or implied of statutory including, without limitation, those of merchantability, fitness for a particular purpose and noninfringement, or arising from a course of dealing, usage, or trade practice.

Table of Contents

Introduction	4
Scope	5
Virtual Chassis Technology Concepts	5
Interconnecting Members in a Virtual Chassis Configuration	5
Virtual Chassis Ports	5
Virtual Chassis Member Roles	6
Master Role	7
Backup Role	7
Line Card Role	7
Graceful Routing Engine Switchover	7
Mastership Priority Setting	8
Mastership Election Process	8
Member ID Numbering	9
Forwarding Path	10
Software Compatibility	10
Automatic Software Upgrades	10
Design Considerations	10
How Many Members Should a Virtual Chassis Configuration Have?	10
Location of Master and Backup Switches	11
Cabling Options	12
Daisy-Chain Ring	12
Braided-Ring	12
Extended Virtual Chassis Configuration	13
Using Uplinks	13
Link Aggregation	14
Implementation	14
Virtual Chassis Split	14
Dynamic Installation	14
Pre-Provisioning	15
Managing and Maintaining a Virtual Chassis Configuration	16
Adding a Member Switch	16
Adding a New Switch to an Existing Virtual Chassis Configuration in the Same Wiring Closet	16
Adding a New Switch from a Remote Wiring Closet to an Existing Virtual Chassis Configuration	17
Replacing a Member Switch	17
Removing, Repairing, and Reinstalling the Same Switch	17
Removing a Member Switch, Replacing it with a Different Switch, and Reapplying the Old Configuration	18
Removing a Member Switch and Making Its Member ID Available for Reassignment to a Different Switch	18
Assigning an IP Address to a Virtual Chassis Configuration	18
Synchronizing Virtual Chassis Members	18

Upgrading Software	18
Loading Factory Default Settings	19
Monitoring a Virtual Chassis Configuration	19
Monitoring Operation with the J-Web Interface	19
Monitoring Operation with CLI Commands	19
Virtual Chassis Technology Deployment Options	20
Data Center, Top of Rack Deployment	20
Data Center, Multi-Row, Top of Rack Deployment	21
Data Center, End of Row Deployment	22
Campus or Building Wiring Closet Deployment	22
Campus or Building Aggregation Deployment	23
Virtual Chassis Ring Deployment	24
Best Practices	25
Conclusion	26
Appendix: Glossary	26
About Juniper Network	27

Table of Figures

Figure 1: Virtual Chassis configuration	4
Figure 2: Virtual Chassis ports on an EX4200 switch rear panel	5
Figure 3: Virtual Chassis port on an EX4200 switch front panel	6
Figure 4: Member roles in a Virtual Chassis configuration	6
Figure 5: Member ID assignments	9
Figure 6: Recommended member location in a daisy-chained ring deployment	11
Figure 7: Recommended member location in a braided-ring deployment	11
Figure 8: Recommended member location in an extended Virtual Chassis configuration	11
Figure 9: Dedicated Virtual Chassis daisy-chained ring	12
Figure 10: Dedicated Virtual Chassis braided-ring	12
Figure 11: Extended Virtual Chassis configuration	13
Figure 12: Uplinks in a Virtual Chassis configuration	13
Figure 13: Data center, top of rack deployment in a single rack	20
Figure 14: Data center, top of rack deployment, daisy-chained ring cabling configuration (single switch per rack shown for simplification)	21
Figure 15: Data center, top of rack deployment, braided-ring cabling configuration (single switch per rack shown for simplification)	21
Figure 16: Data center, multi-row, top of rack deployment (single switch per rack shown for simplification)	21
Figure 17: Data center, end of row deployment (single switch per rack shown for simplification)	22
Figure 18: Campus or building wiring closet deployment, single wiring closet	22
Figure 19: Campus or building wiring closet deployment, multiple wiring closets	23
Figure 20: Campus aggregation, single Virtual Chassis configuration	23
Figure 21: Virtual Chassis ring inter-building deployment example	24
Figure 22: Virtual Chassis ring intra-building example	25

Introduction

Juniper Networks® Virtual Chassis technology is a feature of the Juniper Networks EX4200 line of Ethernet switches allowing the interconnection and operation of switches as a unified, single, high bandwidth device. Up to 10 EX4200 switches may be interconnected via dedicated Virtual Chassis ports on each device, or through optional uplink module ports that are configured as Virtual Chassis ports, with a combined backplane bandwidth of up to 128 Gbps. All EX4200 switch models support Virtual Chassis technology. With this technology, a single logical device that supports up to 480 10/100/1000BASE-T ports or 240 100BASE-FX/1000BASE-X ports may be configured. Optional Gigabit Ethernet or 10-Gigabit Ethernet uplink ports can extend the Virtual Chassis configuration over greater distances. Different EX4200 switch models may be interconnected in the same Virtual Chassis configuration, providing flexibility of port and density options.

Solutions that use the EX4200 line with Virtual Chassis technology combine the scalability and compact form factor of standalone switches with the high availability, high backplane bandwidth characteristics, and high port densities of traditional chassis-based switches. Virtual Chassis configurations enable economical deployments of switches that deliver network availability in locations where installation might otherwise be cost prohibitive or physically impossible.

In a Virtual Chassis configuration, all member switches are managed and monitored as a single logical device. This approach simplifies network operations, allows the separation of placement and logical groupings of physical devices, and provides efficient use of resources. The Virtual Chassis solution offers the same Routing Engine redundancy features as other Juniper Networks chassis-based switches and routers, including graceful Routing Engine switchover (GRES) for hitless failover.

For resiliency and redundancy, the Virtual Chassis configuration includes a master and a backup switch, both dynamically elected as part of the Virtual Chassis configuration. Each remaining switch serves as a line card switch but is ready to be selected as a backup switch if the master or backup switch fails. Switches may also be selectively prioritized in a Virtual Chassis configuration to assign master and backup roles, and to determine the order in which the remaining switches are elected if the master and backup switches fail.

Management of the Virtual Chassis configuration is performed through the master switch. A Virtual Management Ethernet (VME) interface allows remote management by connecting to the out-of-band management port of any member switch through a single IP address.

In addition, the Virtual Chassis configuration uses a single Juniper Networks Junos® operating system image file and a single configuration file. The Junos OS of all member switches in a Virtual Chassis configuration can be upgraded simultaneously from the master switch with a single command.

Key benefits of the Virtual Chassis technology include:

- Simplified overall system maintenance and management through a single management interface.
- Pay-as-you-grow scalability, from 24 to 480 10/100/1000 Mbps ports.
- Extension of the Virtual Chassis configuration by up to 50 km with redundant fiber links between physical devices.
- Redundant, internal hot-swappable power supplies.
- Hot-swappable fan tray with redundant blowers.
- Consistent modular Junos OS control plane feature implementation.
- Dual routing engines with graceful Routing Engine switchover (GRES).

Figure 1 illustrates a typical Virtual Chassis configuration using five EX4200 switches.



Figure 1: Virtual Chassis configuration

Scope

This document provides information about the Juniper Networks EX Series Ethernet Switches Virtual Chassis technology. It describes this technology, explains key concepts, and provides information about designing, operating, maintaining, and managing a Virtual Chassis configuration. This document is intended for system engineers, system administrators, and others who have the technical background that is required to understand network operations, network hardware, and software components.

Virtual Chassis Technology Concepts

The following sections provide an overview of Virtual Chassis technology concepts. This information is intended to help you understand how a Virtual Chassis configuration operates in a variety of situations.

Interconnecting Members in a Virtual Chassis Configuration

A Virtual Chassis configuration consists of between two and 10 EX4200 switches connected together through Virtual Chassis ports. Configurations can range from a deployment in which each switch is physically collocated, to a deployment in which switches are separated by up to 50 km and connected via redundant fiber optics. Each switch within a Virtual Chassis configuration is a member of that Virtual Chassis configuration.

A Virtual Chassis configuration consisting of adjacent switches interconnected with special 0.5 meter, 1.5 meter, 3 meter or 5 meter Virtual Chassis port cables is called a dedicated configuration. A Virtual Chassis configuration interconnected via Gigabit Ethernet or 10-Gigabit Ethernet uplink ports is called an extended configuration.

A variety of cabling methods are used to interconnect Virtual Chassis members in a dedicated configuration. The daisy-chained ring cabling method supports a Virtual Chassis configuration with a height or width of up to 5 meters. The braided-ring cabling method supports a Virtual Chassis configuration with a height or width of up to 22.5 meters. (For more detailed information about the daisy-chained ring and the braided-ring cabling methods, please see the “Cabling Options” section.)

For extended configurations, Virtual Chassis members or dedicated Virtual Chassis configurations are interconnected via optional Gigabit Ethernet or 10-Gigabit Ethernet uplink modules, and the fixed Gigabit Ethernet small form-factor pluggable (SFP) transceiver base ports on the EX4200-24F.

A Virtual Chassis configuration can also be formed with a combination of both dedicated Virtual Chassis connections and extended Virtual Chassis connections.

Virtual Chassis Ports

The rear panel of each EX4200 switch includes two dedicated Virtual Chassis ports, designated VCP 0 and VCP 1. Each of these 64 Gbps ports is used exclusively to interconnect EX4200 switches in a Virtual Chassis configuration. The interfaces for these dedicated ports are operational by default when the ports are cabled with dedicated Virtual Chassis port cables. Virtual Chassis ports do not have cabling dependencies, so, for example,

VCP 0 may be interconnected to VCP 0 or VCP 1 of another switch.

Figure 2 illustrates the Virtual Chassis ports on the EX4200 switch rear panel.




Figure 2: Virtual Chassis ports on an EX4200 switch rear panel

EX4200 switches may be interconnected in an extended Virtual Chassis configuration across distances of up to 50 km with redundant fiber links by installing an optional EX-UM-2XFP dual 10-Gigabit Ethernet small form-factor pluggable (XFP) transceiver, an EX-UM-4SFP (four Gigabit Ethernet SFP) uplink module or an EX-UM-2X4SFP (two-port 10-Gigabit Ethernet SFP+/four-port Gigabit Ethernet SFP) in each EX4200 switch member, or by simply using the fixed SFP base ports in the EX4200-24F. These ports are then configured to function as Virtual Chassis ports so that the interconnected switches are recognized as members of the same Virtual Chassis configuration. Multiple uplinks may be used to interconnect extended Virtual Chassis configurations for increased bandwidth and path redundancy.

With Junos OS 9.6, up to four GbE uplink module ports, two 10GbE uplink module ports or eight SFP switch ports on an EX4200-24F configured as Virtual Chassis ports can be combined to form a single point-to-point connection called a link aggregation (LAG) group or bundle. A LAG provides more bandwidth than a single link; creating a LAG over Virtual Chassis ports creates a bigger “pipe” for forwarding traffic between Virtual Chassis member switches connected by those ports, as well as greater operational redundancy.

Note: Beginning with Junos OS 9.3, the 24 fixed Gigabit Ethernet SFP base ports in the EX4200-24F switch can be configured as Virtual Chassis ports to extend Virtual Chassis configurations.

Figure 3 shows the uplink module in the front panel of an EX4200 switch.




Figure 3: Virtual Chassis port on an EX4200 switch front panel

Each switch (non Virtual Chassis port) port on a Virtual Chassis member is numbered x/y/z, where:

- x is the member ID of the member. (Please see the “Member ID Numbering” section for related information.)
- y is the port interface controller (PIC) ID. Uplink module ports are always on PIC 1 and fixed ports are always on PIC 0.
- z is the port number on the uplink or fixed PIC.

For example, port number 0/1/3 indicates the fourth port (3) on the uplink module (1) on the first member (0) in a Virtual Chassis configuration (please see Figure 3).

Virtual Chassis Member Roles

Each member in a Virtual Chassis configuration is assigned a specific role. A role determines the functions the member performs in the configuration.

One member is assigned the master role and is responsible for managing other members in the Virtual Chassis configuration. One member is assigned the backup role and takes over the master role if the master switch fails. All other members are assigned the line card role. The system executes a mastership election algorithm to determine member roles. For more detailed information, please see the “Mastership Election Process” section.

Figure 4 shows an example of member roles in a Virtual Chassis configuration. The sections that follow describe each role in detail.




Figure 4: Member roles in a Virtual Chassis configuration

Master Role

The member switch that operates in the master role in a Virtual Chassis configuration has full control over the Virtual Chassis configuration. It performs the following functions:

- Serves as the active Routing Engine for the Virtual Chassis configuration.
- Manages the member switches in the Virtual Chassis configuration.
- Runs Junos OS for the Virtual Chassis configuration.
- Runs the chassis management processes and network control protocols.
- Calculates and maintains the forwarding table and distributes it to the local CPU, and then to Packet Forwarding Engines (PFEs) in all member switches.
- Receives and transmits routing information.
- Represents all member switches (the hostname and other properties that are assigned to the master switch during setup apply to all members of the Virtual Chassis configuration).
- Holds the active and master copy of the entire Virtual Chassis configuration. Copies of the active configuration can also be synchronized to all member switches by using the commit sync command-line interface (CLI) command.

Backup Role

The member switch that functions in the backup role in a Virtual Chassis configuration performs the following functions:

- Serves as the backup Routing Engine for the Virtual Chassis configuration.
- Maintains synchronization with the master switch so that it can take over the master role if the master fails.
- Runs Junos OS for EX Series switches in a backup role.
- Synchronizes with the master switch protocol states, forwarding table, and other configurations so that it is prepared to preserve forwarding information and maintain network connectivity with no or minimal disruption if the master switch becomes unavailable.

Line Card Role

Each member that functions in the line card role in a Virtual Chassis configuration performs the following functions:

- Runs Junos OS for EX Series switches in a line card role.
- Detects switch error conditions, such as an unplugged cable, on any interfaces that have been configured on it through the master switch and relays this information to the master switch.
- Receives updates about forwarding information from the master switch and programs these updates into the local PFE.

A Virtual Chassis member in a line card role does not run full network control protocols while in that role. However, if a master or backup switch fails, one of the line card switches takes over the backup role (based on the election process that the "Mastership Election Process" section describes).

A Virtual Chassis configuration must have at least three member switches for there to be a line card switch.

Graceful Routing Engine Switchover

Graceful Routing Engine switchover (GRES) may be configured in a Virtual Chassis configuration. GRES allows the kernel states and forwarding state information from the master member to be automatically maintained and copied to the backup switch so that minimal interruption to network communications occurs if the master goes offline. When GRES is configured, the backup switch automatically synchronizes with the master switch to preserve kernel state and forwarding state information. Any kernel updates to the master switch replicate to the backup switch when they occur. If the kernel on the master switch stops operating, or the master switch experiences a hardware failure, or a manual switchover occurs, the backup switch gracefully assumes the master role.

When GRES is disabled and the backup switch assumes the master role in a redundant failover configuration, the PFEs reinitialize to boot-up state before they connect to the new master switch. In a graceful switchover configuration, the PFEs do not reinitialize their state. Instead, they resynchronize their states with the new master switch to minimize interruption to data traffic.

Mastership Priority Setting

In a configuration that is not pre-provisioned, the role (master, backup, or line card) which a member switch performs within the Virtual Chassis configuration can be designated by configuring its mastership priority from 1 to 255. (For information about a pre-provisioned configuration, please see the “Pre-Provisioning” section.)

When an EX4200 switch powers on, it receives the default mastership priority value of 128. When a standalone EX4200 switch is connected to an existing Virtual Chassis configuration (which implicitly includes its own master), we recommend that you explicitly configure the mastership priority of the members that you want to function as the master and backup switches. (For instructions, please see the “Mastership Priority Setting” section.)

We also recommend the following guidelines for assigning mastership priority:

- Specify the same mastership priority value for the master and backup switches in a Virtual Chassis configuration. Doing so helps to ensure a smooth transition from master to backup if the master switch becomes unavailable. This configuration also prevents the original master switch from retaking control from the backup switch when the original master switch comes back online, a situation sometimes referred to as flapping or pre-emption that can reduce the efficiency of system operation.
- Configure the highest possible mastership priority value (255) for the master and backup switches. This configuration ensures that these members continue to function as the master and backup switches when new members are added to the Virtual Chassis configuration.

Mastership Election Process

When a Virtual Chassis configuration boots, the Junos OS for EX Series switches automatically runs the mastership election process to determine which member switches take the role of master, backup, and line cards. All Virtual Chassis member switches participate in the election process. If a master switch fails, the backup switch automatically and immediately takes on the master role, which minimizes interruption to the operation of Virtual Chassis configuration. The system subsequently runs the mastership election process to elect one of the line card switches as the new backup switch. (The system also runs this process if the backup switch fails.)

The election algorithm follows the sequence below to assign member roles and elect a master and a backup switch. The master role is assigned to the switch with the highest ranking. The backup role is assigned to the switch with the second highest ranking. Other switches become line cards.

1. Choose the member with the highest user-configured mastership priority. (One is the lowest and 255 is the highest possible value; please see the “Mastership Priority Setting” section for more information.)
2. Choose the member that was the master switch the last time the Virtual Chassis configuration booted.
3. In a Virtual Chassis configuration merge scenario, choose the (master) member that has the highest number of current members in the existing Virtual Chassis configuration. (A merge scenario occurs when two active Virtual Chassis configurations, each with its own master, are combined.)
4. Choose the member that has been included in the Virtual Chassis configuration for the longest time. (For this factor to be considered, there must be a lapse of at least one minute between the power-on of each interconnected member switch.)
5. Choose the member with the lowest Media Access Control (MAC) address.

If you want to ensure that a specific member is elected as the master switch during initial Virtual Chassis installation, follow these steps:

1. Power on only the switch that you want to configure as master switch of the Virtual Chassis configuration.
2. Configure the mastership priority of that member to have the highest possible value (255).
3. Without connecting to the master switch, power on all other member switches individually and reset their configurations to factory-default values either through the front LCD menu or by using the following configuration mode CLI command:

```
user@host# load factory-default
```

4. Connect all other members to the master switch while they are powered off.
5. Power on the other member switches.
6. Continue to configure other member switches through the master switch, as desired. These member switches default to a mastership priority of 128 if no previous configuration exists.

Member ID Numbering

Each EX4200 switch is a potential member of a Virtual Chassis configuration. When an EX4200 switch powers on, it receives a member ID (which is displayed on its front-panel LCD). If the switch powers on as a standalone switch, its member ID is 0. When the switch interconnects with other EX4200 switches in a Virtual Chassis configuration, the master switch assigns a member ID (0 through 9) to the switch. The member ID is based on a variety of factors, including the order in which the switch was added to the Virtual Chassis configuration. As each switch is added and powered on, it receives the next lowest available (unused) member ID.

The member ID distinguishes the member switches from each other. You use the member ID to:

- Assign a mastership priority value to a member switch.
- Configure interfaces for a member switch (a function that is similar to a slot number on Juniper Networks chassis-based routers and switches).
- Apply various operational commands to a member switch.
- Display status or characteristics of a member switch.

Note: Members can be physically located in any order in a Virtual Chassis configuration. They do not need to be placed in order of member ID.

Figure 5 illustrates member ID assignments in a Virtual Chassis configuration with five members.




Figure 5: Member ID assignments

Member ID assignments follow these guidelines:

- When a Virtual Chassis member reboots, it retains its member ID.
- The member ID of a switch that is removed from a Virtual Chassis configuration is not released automatically to the available member ID pool.
- If the Virtual Chassis configuration previously included a member switch and that member was physically disconnected or removed from the Virtual Chassis, its member ID is not automatically available for re-assignment by the master switch. For example, consider a Virtual Chassis configuration that is composed of member 0, member 2, and member 3 because member 1 was removed. When you add another member switch and power it on, the master switch assigns it a member ID of 4.
- A replacement switch is treated as a new addition to the Virtual Chassis configuration and receives the next lowest available member ID.
- You can use the following operational mode CLI command to manually configure the member ID of a device. Replace <current-member-id> with the existing member ID of the device, and replace <new-member-id> with the member ID that you want to assign to the device.


```
user@host> request virtual-chassis renumber member-id <current-member-id> new-member-id <new-member-id>
```
- You can use the following CLI command to return a member ID that was previously used but is no longer assigned to any active member to the member ID pool. Replace <current-member-id> with the desired member ID.


```
user@host> request virtual-chassis recycle member-id <current-member-id>
```

Forwarding Path

A Virtual Chassis configuration uses an internal shortest path forwarding algorithm to determine the path for routing packets internally through the member switches. When you deploy a Virtual Chassis configuration, its Virtual Chassis Control Protocol (VCCP) builds a forwarding table that includes information about each switch component and its location. From this table, the system determines the shortest forwarding path for data between the ingress port and the egress port in a Virtual Chassis configuration.

Software Compatibility

A Virtual Chassis configuration can include EX4200 switches with different hardware versions and model numbers. However, each member switch must run the same software version as the master switch.

When a member switch is added to a Virtual Chassis configuration, the master switch checks the compatibility of the Junos OS running on the new switch. A switch running a different software version than the master switch obtains a member ID from the master switch, but it does not become a functional member of the Virtual Chassis configuration and does not forward data packets.

The required software package should be downloaded to the master switch in the Virtual Chassis configuration. Then the new member switch can be upgraded to the stored software version using the following CLI command. This command downloads the image from the master switch through the Virtual Chassis ports to the specified member and then reboots the member. The member does not need to be directly connected to the master switch.

```
request system software add <package location> member <member-id> reboot
```

Automatic Software Upgrades

A feature available with Junos OS release 10.0 automatically upgrades the software version on new member switches joining a Virtual Chassis configuration to the same Junos OS version running on the Master switch. This feature is disabled by default; the user must explicitly enable it and specify a path—either to a local directory or a remote server location—where new Virtual Chassis switch members can obtain the appropriate software image.

```
set virtual-chassis auto-sw-update package-name <image-path>
```

Design Considerations

The following sections describe design considerations of a Virtual Chassis configuration. These sections include information about cabling Virtual Chassis members together, appropriate number of members, and link aggregation across Virtual Chassis members.

How Many Members Should a Virtual Chassis Configuration Have?

A Virtual Chassis configuration may include between two and ten switches. When selecting the appropriate number of switches for a deployment, there are a number of factors to consider, including port density, resiliency requirements, and system cost.

In general, a Virtual Chassis configuration in which ports are distributed across as many switches as possible provides the highest availability. However, increasing the number of switches also increases cost and space requirements. For example, assume that you want to deploy a system that includes 96 ports. Options for this system include the following:

- Option 1:
 - Two EX4200-48x switches—one switch serving in the master role and one switch serving in the backup role.
 - Advantage: Compact footprint, cost-effective.
 - Disadvantage: Loss of one switch affects 50 percent of users.
- Option 2:
 - Four EX4200-24x switches—one switch serving in the master role, one switch serving in the backup role, and two switches serving as line cards.
 - Advantage: Higher availability as the loss of one switch only affects 25 percent of users; does not affect uplinks if the failed switch did not include any uplinks.
 - Disadvantage: Increased space, power, and cost.

Location of Master and Backup Switches

When selecting the physical placement of members in a Virtual Chassis configuration, consider the following guidelines. These guidelines are intended to ensure the lowest probability of a severe fault bringing down both the master and backup switches. The guidelines also ensure a high probability that, if a Virtual Chassis configuration splits, the master and backup switches fall into different parts, and each part has at least half the original Virtual Chassis members. (A Virtual Chassis split is an unusual situation that may occur as the result of a double failure in a Virtual Chassis configuration ring topology.)

- Master and backup switches should be evenly spaced by member hop in a Virtual Chassis configuration.
- In an extended Virtual Chassis configuration, place the master and backup switches in separate locations. (Please see Figure 8 for an example, and note the multiple uplink ports configured as Virtual Chassis ports for added bandwidth and link redundancy.)

Figure 6 illustrates the recommended location of the master and backup switches in a daisy-chained ring cabling configuration.




Figure 6: Recommended member location in a daisy-chained ring deployment

Figure 7 illustrates the recommended location of the master and backup switches in a braided-ring cabling configuration.




Figure 7: Recommended member location in a braided-ring deployment

Figure 8 illustrates the recommended location of the master and backup switches in an extended Virtual Chassis configuration. The top switch is connected to the bottom switch in each stack for added bandwidth and redundancy.




Figure 8: Recommended member location in an extended Virtual Chassis configuration

Cabling Options

Physical placement of EX4200 switches as part of a Virtual Chassis configuration is flexible. Possible deployments include members in a single rack, across several racks, in the same wiring closet, and spanning wiring closets across floors, buildings, and facilities. When interconnecting switches through dedicated Virtual Chassis ports, the physical distance between two directly connected switches may not exceed 5 meters, which is the maximum Virtual Chassis port cable length. A Virtual Chassis configuration may be extended by using uplink ports configured as Virtual Chassis ports to allow a greater distance between two directly connected member switches.

There are three cabling methods for interconnecting switches in a Virtual Chassis configuration: daisy-chained ring, braided ring, and extended Virtual Chassis configuration. The following sections describe each method. We recommend that switches in a Virtual Chassis configuration be connected in a ring topology for resiliency and speed. A ring configuration provides up to 128 Gbps of bandwidth between member switches.

Daisy-Chained Ring

In the daisy-chained ring configuration, each device in the Virtual Chassis configuration is connected to the device immediately adjacent to it. In addition, members at the end of the Virtual Chassis configuration are connected to each other to complete the ring topology. Connections between devices can use either Virtual Chassis port on the back of a device (for example, VCP 0 to VCP 0 or VCP 0 to VCP 1).

The daisy-chained ring configuration provides a simple and intuitive method for interconnecting devices. The maximum height or breadth of the Virtual Chassis is 5 meters.

Figure 9 illustrates the daisy-chained ring cabling method.




Figure 9: Dedicated Virtual Chassis daisy-chained ring

Braided-Ring

In the braided-ring cabling configuration, alternating devices in the Virtual Chassis configuration are connected to each other. In addition, the two device pairs at each end of the Virtual Chassis configuration are directly connected to each other to complete the ring topology. Connections between devices can use either Virtual Chassis port on the back of a device.

The braided ring configuration extends the Virtual Chassis height or breadth to 22.5 meters.

Figure 10 illustrates the braided-ring cabling method.




Figure 10: Dedicated Virtual Chassis braided-ring

Extended Virtual Chassis Configuration

The extended Virtual Chassis configuration allows the interconnection of individual Virtual Chassis members or dedicated Virtual Chassis configurations across distances of up to 50 km with redundant fiber links. This configuration is used when deploying a Virtual Chassis configuration across wiring closets, data center racks, data center rows, or facilities. In this configuration, optional EX-UM-2XFP or EX-UM-4SFP uplink modules or fixed SFP base ports in the EX4200-24F are used to interconnect the members of the Virtual Chassis configuration. Multiple uplinks may be used for additional bandwidth and link redundancy.

To implement this configuration, use the following CLI command:

```
user@host> request virtual-chassis vc-port set pic-slot <pic-slot> port <port> member <member-id>
```

Note: Beginning with Junos OS 9.3, the 24 fixed Gigabit Ethernet SFP base ports in the EX4200-24F switch can be configured as Virtual Chassis ports to extend Virtual Chassis configurations.

Figure 11 illustrates the extended Virtual Chassis configuration.




Figure 11: Extended Virtual Chassis configuration

Using Uplinks

Each EX4200 switch supports a single optional uplink module. Two 10-Gigabit Ethernet small form-factor pluggable transceivers (XFPs) or four Gigabit Ethernet SFP uplink ports are supported per uplink module and may be used to interconnect other devices, such as another switch when configured as either a Virtual Chassis port (that is, the connected switch must be another EX4200 switch configured with uplink Virtual Chassis ports) or as a standard Layer 2/3 uplink.

When using uplinks, we recommend that:

- Uplink modules be placed in a Virtual Chassis configuration line card switches. This approach prevents the loss of a master or backup member and an uplink port if a single device fails.
- Uplinks be placed in devices separated at equal distance by member hop. This approach ensures the best Virtual Chassis configuration uplink-to-backplane traffic distribution and the lowest probability of oversubscribing the backplane. It also provides the highest uplink coverage if a member switch or Virtual Chassis port fails or a Virtual Chassis configuration split occurs.

Figure 12 illustrates the use of uplinks in a Virtual Chassis configuration.

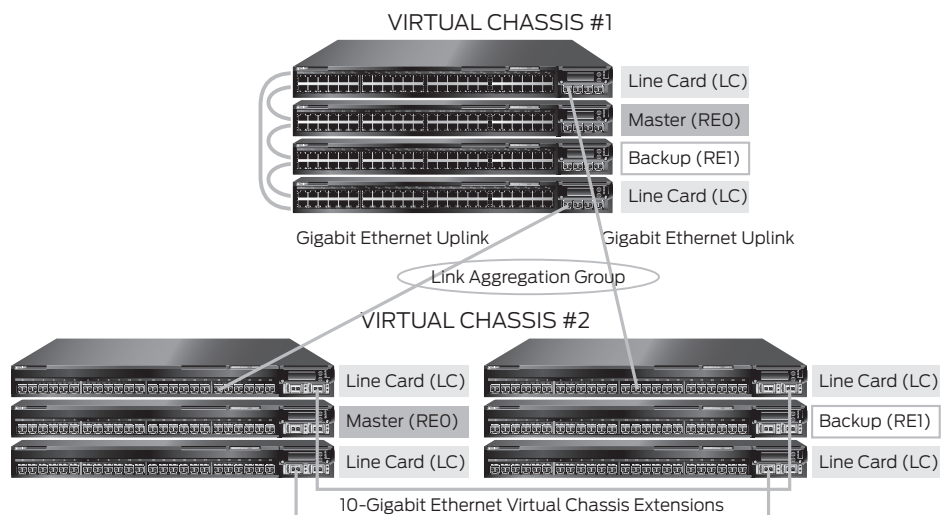


Figure 12: Uplinks in a Virtual Chassis configuration

Link Aggregation

Multiple physical switch ports may be grouped into a logical point-to-point link, known as a link aggregation group (LAG). A LAG provides increased bandwidth and link redundancy. If one link in the LAG fails, the system automatically redirects traffic across remaining links.

Up to eight ports may be included within a LAG in a Virtual Chassis configuration. Ports may be distributed across the fixed or uplink ports of any members. A Virtual Chassis configuration can support up to 64 LAGs.

Implementation

The following sections describe options for installing a Virtual Chassis configuration. There are two methods for installing and configuring Virtual Chassis technology: dynamic and pre-provisioning.

The dynamic method offers a simple plug-and-play option for building a Virtual Chassis configuration. It does not require any configuration. However, this method does not allow user selection of the master and backup switches, and it does not prevent certain user errors, such as adding the wrong switch into a Virtual Chassis configuration.

The pre-provisioning method requires careful planning and user configuration before installing the Virtual Chassis configuration. However, because all member switches and their roles in a given Virtual Chassis must be statically defined, this method minimizes user error and provides consistent and deterministic results if a member switch fails.

Virtual Chassis Split

In rare cases, Virtual Chassis configurations can be “split” when a single switch, cable or Virtual Chassis port in a chained topology fails, or when two or more adjacent switch members lose connectivity in a ring topology. Prior to Junos OS 9.3, such a split would initiate a new Master re-election process in each of the new segments, resulting in two or more identical, active Virtual Chassis configurations unaware of each other’s presence. This could create problems at the network level when, for example, each of the new Virtual Chassis segments contain the same routing information, resulting in two or more “routers” in the network with identical IDs, networks, advertisements, etc.

To alleviate this problem, beginning with Junos OS 9.3, an enhancement option called Virtual Chassis “split-detection” has been added so that, following a Virtual Chassis split, no more than one segment would remain active and operational under the original configuration. All other segments created by the split are forced into an “inactive” Virtual Chassis state, with all switch members assuming a line card role with no Master or Backup. Users must either manually reload the factory default configuration on the inactive switch members or resolve the faults that caused the split (e.g. repair or replace the broken Virtual Chassis cables or member switches) before the switches can rejoin the active Virtual Chassis configuration. Beginning with Junos OS 9.3, Virtual Chassis split detection is enabled by default; for more details about Virtual Chassis split detection and split rules, please refer to the appropriate Juniper EX4200 switch product manual.

While the Virtual Chassis split detection implementation is desirable in most cases, there are rare situations that could lead to all parts of a Virtual Chassis split being rendered inactive. For example, if a split occurs and the resulting segments have an equal number of Virtual Chassis members, the original Master and Backup switches do not fall into the same split segment, and the original Backup switch is part of the failure that triggers the split, the resulting split parts would have no “active” members. A real-world example of such a situation would be a two-member Virtual Chassis configuration. If the backup switch fails for some reason, the remaining Master switch would interpret this as a Virtual Chassis split and would put itself into an “inactive” state based on the Virtual Chassis split rules, bringing the entire Virtual Chassis down. In such cases, the Virtual Chassis split-detection option would not be desirable and should be disabled with the following CLI command:

```
set virtual-chassis no-split-detection
```

Dynamic Installation

Use the dynamic installation method to build a Virtual Chassis configuration or to add a new member to an existing Virtual Chassis configuration without prior user configuration. Although it is not required, we recommend that the master and backup switches be designated by configuring the mastership priority of these switches to be the highest value of all members.

We also recommend that factory defaults be loaded on all members before adding these switches to the Virtual Chassis configuration. This procedure prevents unexpected behavior during the addition of the new member. Factory defaults can be loaded in either of these ways:

- Using the following configuration mode CLI commands:

```
user@host# load factory-default
user@host# set system root-authentication plain-password
```

(Follow the prompts to configure a root password to apply the change.)

```
user@ host#commit
```
- Using the LCD menus on a switch

To perform a dynamic installation, follow these general steps:

1. Begin by installing the switch that you want to be the master by taking these actions:
 - a. Install required power supplies.
 - b. Place the switch in the desired location.
 - c. Power up the switch and load the factory default configuration.
 - d. Assign the highest mastership priority to the switch.

This switch becomes the master and obtains a member ID of 0.

2. For each switch that you want to be a line card switch, take these actions:
 - a. Install required power supplies.
 - b. Power up the switch and load the factory default configuration.
 - c. Power off the switch.
 - d. Place the switch in the desired location.
 - e. Connect the switch to existing members of the Virtual Chassis with a Virtual Chassis port cable. (For instructions, please see the “Cabling Options” section.)
 - f. Power up the switch.
 - g. Repeat these steps as needed to install additional line card switches.

The mastership election process assigns the switch the next lowest available member ID based on the order in which the switch is added to the Virtual Chassis configuration. The switch may be temporarily configured as the backup switch (because the proposed backup switch has not yet been added to the Virtual Chassis configuration). The switch becomes a line card when the actual backup switch is installed.

3. For the switch that you want to be the backup switch in the Virtual Chassis configuration:
 - a. Install required power supplies.
 - b. Power up the switch and load the factory default configuration.
 - c. Power off the switch.
 - d. Place the switch in the desired location.
 - e. Connect the switch to existing members of the Virtual Chassis configuration with a Virtual Chassis port cable. (For instructions, please see the “Cabling Options” section.)
 - f. Power up the switch.

The mastership election process assigns the switch the next lowest available member ID based on the order that the switch is added to the Virtual Chassis configuration.

- g. Assign this switch the same mastership priority as the master switch. Doing so prevents mastership preemption if the master fails and then recovers.

Pre-Provisioning

A pre-provisioned configuration allows you to deterministically control the member ID and role assigned to a member switch by associating the switch to its serial number. A pre-provisioned configuration file links the serial number of each EX4200 switch to a designated member ID and role. The serial number must be specified in the configuration file for the member to be recognized as part of the Virtual Chassis configuration.

In this configuration, you must select two members that you want to make eligible for election as the master and backup switches. When you list these two members in the pre-provisioned configuration, you designate the member role as Routing Engine. One member then functions as the master switch of the Virtual Chassis configuration and the other functions as the backup switch.

Additional members, which are not eligible for election as the master or backup switch, can be specified as line cards in the pre-provisioned configuration.

The pre-provisioned configuration provides the option of not explicitly assigning a role to a member switch. This member switch is eligible for election as the backup switch if the master or the backup switch fails. It can also become the master switch if both the master and backup switches fail.

In a pre-provisioned configuration:

- A switch member can be explicitly configured as a routing engine and can then become either master or backup, depending on the results of the master election process.
- You can explicitly configure a member with the role of line card, which makes it ineligible for functioning as a master or backup switch.
- A member that is not explicitly assigned a role is eligible to become either the backup switch (if the master switch or backup switch fails) or the master switch (if both the master switch and backup switches fail).
- The mastership priority value is assigned by the switch software based on the specified role:
 - The master and backup switches are assigned a mastership priority of 129.
 - A line card switch is assigned a mastership priority of 0, making it ineligible to participate in the master election.
 - A switch that is not explicitly assigned a role is configured with a mastership priority of 128, making it eligible to participate in the master election.

Managing and Maintaining a Virtual Chassis Configuration

The following sections provide information about adding or replacing member switches, and they provide guidance for a variety of other maintenance operations.

Note: Some maintenance operations require connecting to the Virtual Chassis configuration through a console so that you can issue CLI commands. If you connect through a console to a Virtual Chassis member other than the master switch, the connection redirects to the master switch automatically.

Adding a Member Switch

The following sections provide information about adding a member switch to a Virtual Chassis configuration.

Adding a New Switch to an Existing Virtual Chassis Configuration in the Same Wiring Closet

This section describes how to add a new switch to a Virtual Chassis configuration when the new switch is to be in the same wiring closet as the other member or members of the Virtual Chassis configuration. Before beginning:

- Install required power supplies and uplink modules in the new switch.
- Place the new switch in the desired location.
- Confirm that the new switch is powered off.
- If expanding a pre-provisioned configuration, make a note of the serial number on the back of the switch. Then edit the existing Virtual Chassis pre-provisioning configuration to include the serial number and the role of the new switch.

To add the new switch, follow these steps:

1. If the new member switch was previously configured, power it on, reset the configuration to the factory default values, and then power it off. (For instructions about how to load factory default values, please see the “Dynamic Installation” section.)
2. Interconnect the unpowered new switch to at least one member of the existing Virtual Chassis configuration by using the dedicated Virtual Chassis ports.
3. Power on the new switch.

4. Confirm that the new switch is now included within the Virtual Chassis configuration by checking the front-panel LCD for the member ID. It should display a member ID that is higher than 0 because there are already at least two members of the Virtual Chassis configuration.

If you are using a pre-provisioned configuration, the member ID is assigned to the serial number of the member in the configuration file.

Adding a New Switch from a Remote Wiring Closet to an Existing Virtual Chassis Configuration

To add a new switch from a remote wiring closet to an existing Virtual Chassis configuration, use the EX-UM-2XFP or EX-UM-4SFP uplink ports. Configure the uplink ports on both sides of the link as Virtual Chassis ports.

Before beginning:

- Power on the new switch as a standalone device to configure its uplinks as Virtual Chassis ports.
- Install required power supplies and uplink modules in the new switch.
- Place the new switch in the desired location.
- If the new switch was previously configured, reset the configuration to the factory default values. (For instructions about how to load factory default values, please see the “Dynamic Installation” section.)
- If expanding a pre-provisioned configuration, make a note of the serial number on the back of the switch. Then edit the existing Virtual Chassis configuration to include the serial number of the new member switch. Specify the member ID and the role of the new switch when adding its serial number in the Virtual Chassis pre-provisioned configuration. The parameters which are specified in the master Virtual Chassis configuration file are applied after the new member switch interconnects with its uplink Virtual Chassis port.
- Prepare an existing member for interconnecting with the new switch through an uplink port by configuring an uplink port as a Virtual Chassis port on the existing member.

To add the new switch, follow these steps:

1. Interconnect the new member switch to at least one member of the existing Virtual Chassis configuration, using the uplink ports that have been configured as Virtual Chassis ports on both sides of the uplink.
2. Power on the new switch.
3. Confirm that the new switch is included within the Virtual Chassis configuration by checking the front-panel LCD for the proper member role and member ID.

If using a pre-provisioned configuration, the member ID and member role are assigned to the member serial number in the configuration file.

Replacing a Member Switch

A member switch may be replaced in a Virtual Chassis configuration with little loss of traffic. When a switch is replaced, the configuration of the replaced switch is retained and may be applied to the new switch if desired.

The following sections provide information about replacing a member switch in a variety of situations.

Note: A replacement switch does not inherit the configuration of the previous switch if its member ID changes.

Removing, Repairing, and Reinstalling the Same Switch

If a member switch needs replacement, it may be removed from the Virtual Chassis configuration with little loss of traffic. The master switch stores the configuration, including the member ID of the removed switch, so that it can be reapplied when the switch (with the same MAC address) is reconnected.

To remove, repair, and reinstall a switch, follow these steps:

1. Power off and disconnect the switch.
2. Repair the switch, as necessary.
3. Reconnect and power on the switch.

Removing a Member Switch, Replacing it with a Different Switch, and Reapplying the Old Configuration

A member switch may be replaced with a different switch that retains the configuration of the original switch. The master switch stores the configuration of the member that was removed. When a replacement member switch is connected, the master switch assigns a new member ID to it. The old configuration remains stored under the member ID of the previous member switch. Renumbering the replacement switch with the member ID of the old switch applies the configuration of the old switch to the replacement switch.

To replace a switch and reapply configuration, perform the following steps:

1. Power off and disconnect the switch to be replaced.
2. If the replacement switch has been previously configured, reset its configuration to the factory default values. (For instructions about how to load factory default values, please see the “Dynamic Installation” section.)
3. Connect the replacement switch in place of the old switch and power it on.
4. Note the member ID that displays on the front-panel LCD. It should be the next lowest member ID that was available before the replacement.
5. Use the following CLI command to change the current member ID of the replacement switch to the member ID of the switch that you removed:

```
user@host> request virtual-chassis renumber member-id <member-id> new-member-id <member-id>
```

6. If you used a pre-provisioned configuration, use the following CLI command to change the relevant serial number in the Virtual Chassis configuration file. Substitute the serial number of the replacement switch (on the back of the switch) for the serial number of the switch that was removed.

```
user@host# set virtual-chassis pre-provisioned member <removed member-id> serial-number <replacement switch serial number> role <removed member old role>
```

Removing a Member Switch and Making Its Member ID Available for Reassignment to a Different Switch

When removing a member switch from the Virtual Chassis configuration, the master switch keeps the member ID of that switch on reserve.

In addition, the configuration of the old switch applies only to the valid parts of the new member. For example, if a 24-port switch is replaced with a 48-port switch, the old configuration applies to the first 24 ports of the new switch. The other ports have no configuration.

Assigning an IP Address to a Virtual Chassis Configuration

A Virtual Chassis configuration is managed as a single logical network element. As such, it has only one management IP address, which is configured on the Virtual Management Ethernet (VME) interface. This VME interface is a logical IP interface associated with the Virtual Chassis internal management VLAN that connects the me0 interfaces of all member switches in a Virtual Chassis configuration. To assign an IP address, use these CLI commands:

```
user@host> configure
```

```
user@host# set interfaces vme unit 0 family inet address <ip-address>/<subnet-mask>
```

Synchronizing Virtual Chassis Members

Whenever the configuration settings on the master switch are changed, we recommend propagating changes to all other switches in the Virtual Chassis configuration. To do so, use the following configuration mode CLI command:

```
user@host> configure
```

```
user@host# commit synchronize
```

Upgrading Software

For ease of management, the Virtual Chassis configuration provides flexible methods for upgrading software releases. A software upgrade may be applied to the entire Virtual Chassis configuration or to any member in the Virtual Chassis configuration through a CLI or Juniper Networks J-Web Software command.

To upgrade the software image to all members in a Virtual Chassis configuration, issue the following CLI command on the master switch. The master switch automatically applies the upgrade to all members in the Virtual Chassis configuration and then reboots to complete the upgrade.

```
user@host> request system software add <image-name> reboot
```

Loading Factory Default Settings

Each EX Series switch is programmed with a factory default configuration stored in a default configuration file. This default configuration includes values for system parameters such as Syslog and commit, configures Power over Ethernet (PoE) and Ethernet switching on all interfaces, and enables the Link Layer Discovery Protocol (LLDP) and Rapid Spanning Tree Protocol (RSTP).

When committing changes to the configuration, a new configuration is created and becomes the active configuration. You can revert to the factory default configuration at any time. (For instructions about how to load factory default values, please see the “Dynamic Installation” section.)

Monitoring a Virtual Chassis Configuration

The following sections provide information about monitoring the operation of a Virtual Chassis configuration.

Monitoring Operation with the J-Web Interface

The J-Web interface may be used to view the following information about Virtual Chassis members and ports:

- Member details and how members are connected with each other.
- Traffic statistics between the Virtual Chassis ports of selected members.

To view Virtual Chassis configuration monitoring details in the J-Web interface, select Monitor > Virtual Chassis.

The top part of the screen displays details of the Virtual Chassis configuration, including:

- Member ID
- Priority
- Role
- Interface
- The member ID of the neighboring switch

In the bottom part of the screen, select a member ID to view Virtual Chassis port input and output rates, and select the interval at which the charts must be refreshed. Click the Stop button to stop fetching values from the switch, and click the Start button to start plotting data again from the point at which you stopped.

Monitoring Operation with CLI Commands

The Junos OS CLI provides commands which display information for all members in a Virtual Chassis configuration or for one specific member.

To view member details for all members in a Virtual Chassis configuration, enter the following CLI command:

```
user@host> show virtual-chassis status
```

This command displays the following information:

- Virtual Chassis ID
- Member ID
- Status
- Serial number
- Model
- Membership priority
- Role
- Neighbor list

To view Virtual Chassis port traffic statistics for a specific member in a Virtual Chassis configuration, enter the following CLI command, replacing <member-id> with the ID of the member for which you want to view information:

```
user@host> show virtual-chassis vc-port statistics member <member-id>
```

This command displays the following information:

- Member ID
- Virtual Chassis port information

Virtual Chassis Technology Deployment Options

The following sections provide overviews of various deployments of the Virtual Chassis configuration.

Data Center, Top of Rack Deployment

The data center, top of rack deployment is suitable for data center environments in which members of the Virtual Chassis configuration are colocated with servers in the same rack.

A single Virtual Chassis configuration consisting of multiple switches in the same rack simplifies management by reducing the number of logically managed devices, and it offers flexible options for the number and deployment of uplinks. This deployment also provides servers with the capability of configuring NIC teaming (LAG) to multiple members of the same Virtual Chassis configuration, increasing the total server network bandwidth while providing server link redundancy.

Figure 13 illustrates a data center, top of rack deployment in a single rack.

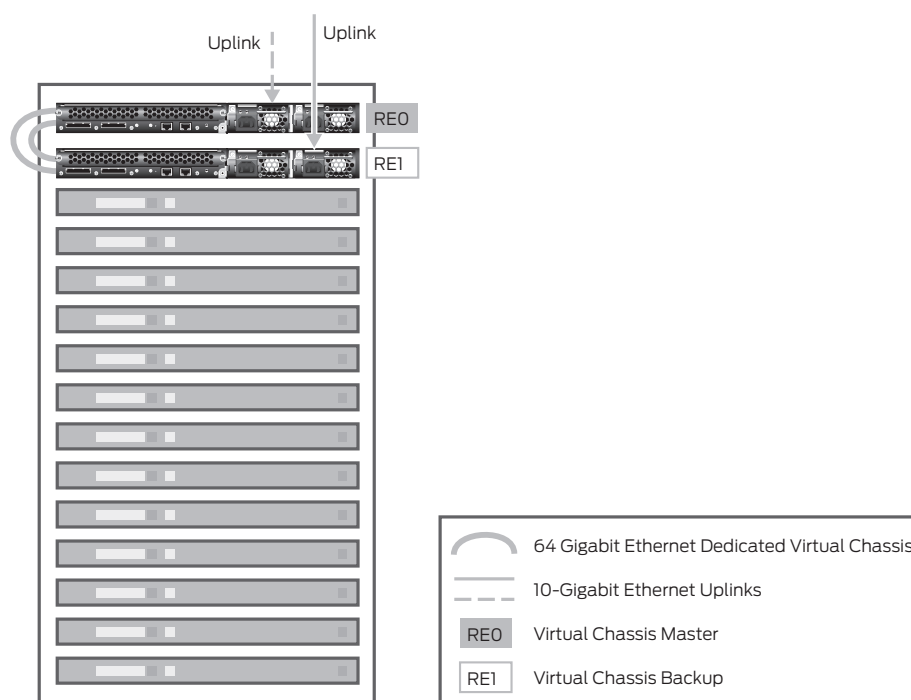


Figure 13: Data center, top of rack deployment in a single rack

The Virtual Chassis configuration may be extended across multiple racks within the same row in a data center. Figure 14 shows an example of extending the data center, top of rack deployment across several racks. This figure illustrates a daisy-chained, ring cabling configuration, which allows spacing of up to 5 meters between adjacent racks and requires one or more uplink extension to complete the ring. Extending the Virtual Chassis across racks further eases management, lowers inter-rack server-to-server latency, and increases uplink flexibility, ultimately leading to lower total cost of ownership (TCO).

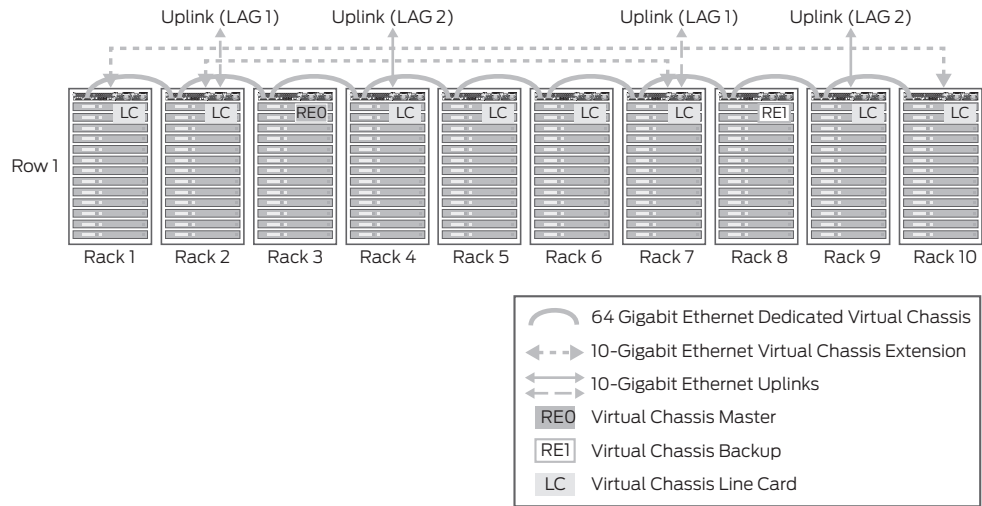


Figure 14: Data center, top of rack deployment, daisy-chained ring cabling configuration (single switch per rack shown for simplification)

Figure 15 shows an example of extending the data center, top of rack deployment across several racks. This figure illustrates a braided-ring cabling configuration, which allows spacing of up to 5 meters between alternate racks.

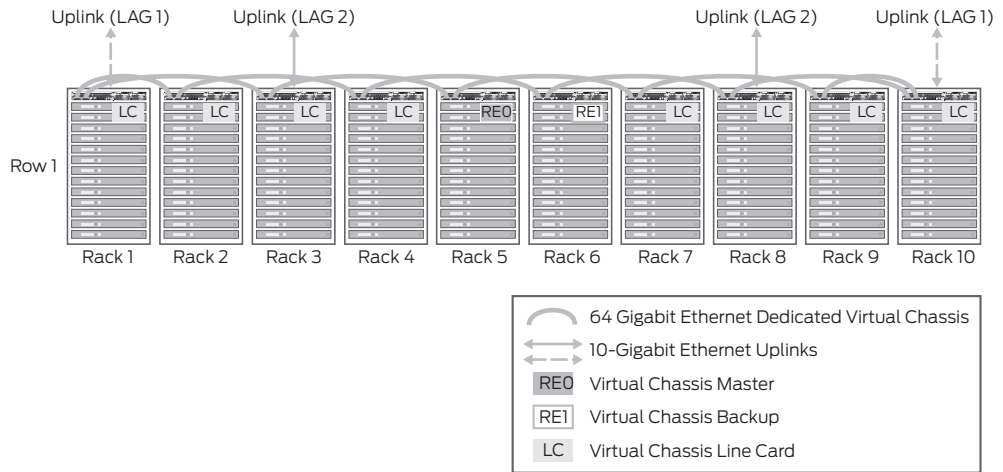


Figure 15: Data center, top of rack deployment, braided-ring cabling configuration (single switch per rack shown for simplification)

Data Center, Multi-Row, Top of Rack Deployment

The Virtual Chassis configuration may be extended across multiple racks and across rows in a data center, as shown in Figure 16.

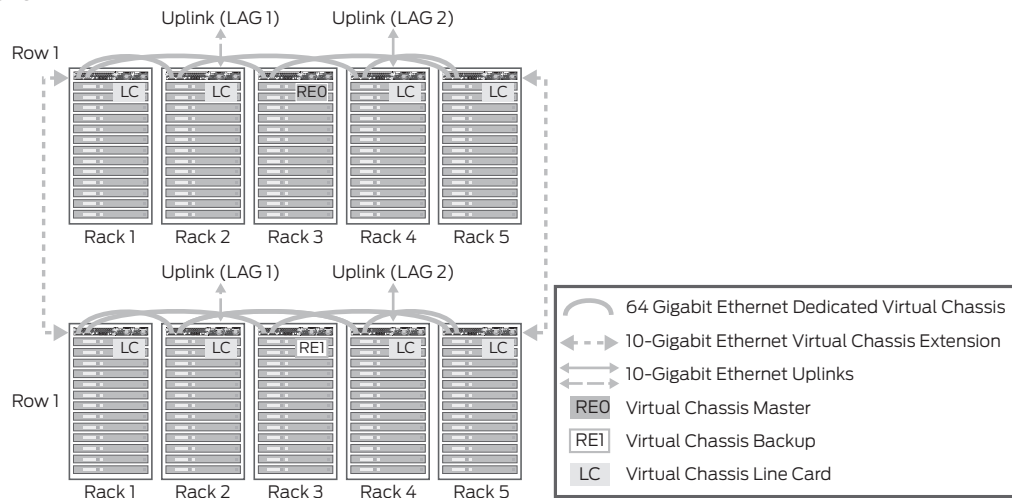


Figure 16: Data center, multi-row, top of rack deployment (single switch per rack shown for simplification)

Data Center, End of Row Deployment

The data center, end of row deployment allows the installation of EX4200 switches as end-of-row devices in a data center. Figure 17 illustrates this deployment. Virtual Chassis members are interconnected by using the daisy-chained or braided-ring cabling configuration, and uplinks provide connectivity to core switching devices.




Figure 17: Data center, end of row deployment (single switch per rack shown for simplification)

Campus or Building Wiring Closet Deployment

The Virtual Chassis configuration may be used in campus or building wiring closets. A single wiring closet may include up to 480 10/100/1000 ports and 20 10-Gigabit Ethernet or 40 Gigabit Ethernet uplinks. The Virtual Chassis configuration may be extended across multiple wiring closets to ease management and provide design flexibility for uplinks into the LAN core. Figure 18 illustrates the deployment of a Virtual Chassis configuration within a single wiring closet. Figure 19 illustrates the deployment of a Virtual Chassis configuration when the Virtual Chassis configuration spans more than one wiring closet.




Figure 18: Campus or building wiring closet deployment, single wiring closet

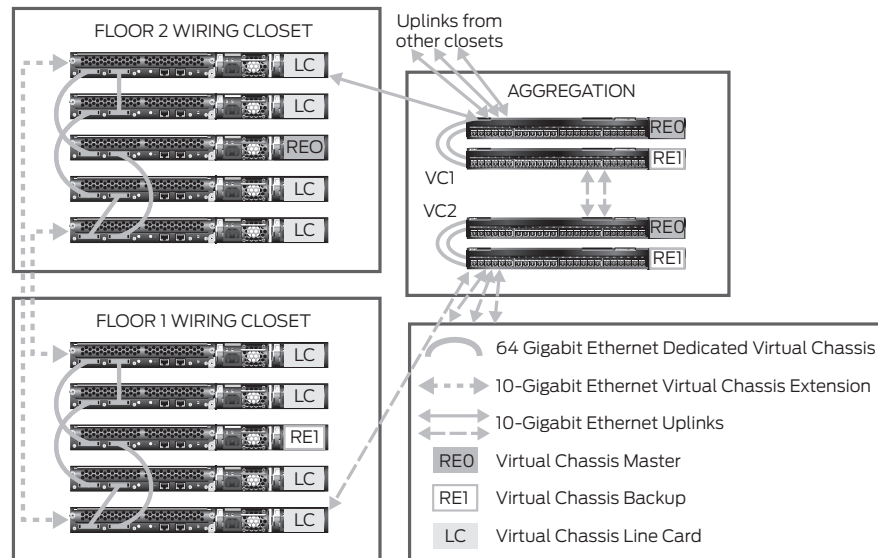


Figure 19: Campus or building wiring closet deployment, multiple wiring closets

Campus or Building Aggregation Deployment

The EX4200-24F (featuring 24 SFP ports) is designed for aggregating access switch devices from wiring closets or data center deployments. It provides options of physical placement and port density to meet a range of requirements. For maximum availability, a pair of Virtual Chassis configurations may be located in the aggregation layer with each access switch (or Virtual Chassis configuration) dual-homed. Or, a single Virtual Chassis configuration may be deployed at the aggregation layer, providing ease of management and eliminating the active use of layer 2 redundancy protocols such as Spanning Tree.

Figures 18 and 19 above illustrate a deployment for a separate Virtual Chassis configuration at aggregation. Figure 20 illustrates a single Virtual Chassis deployment at aggregation.

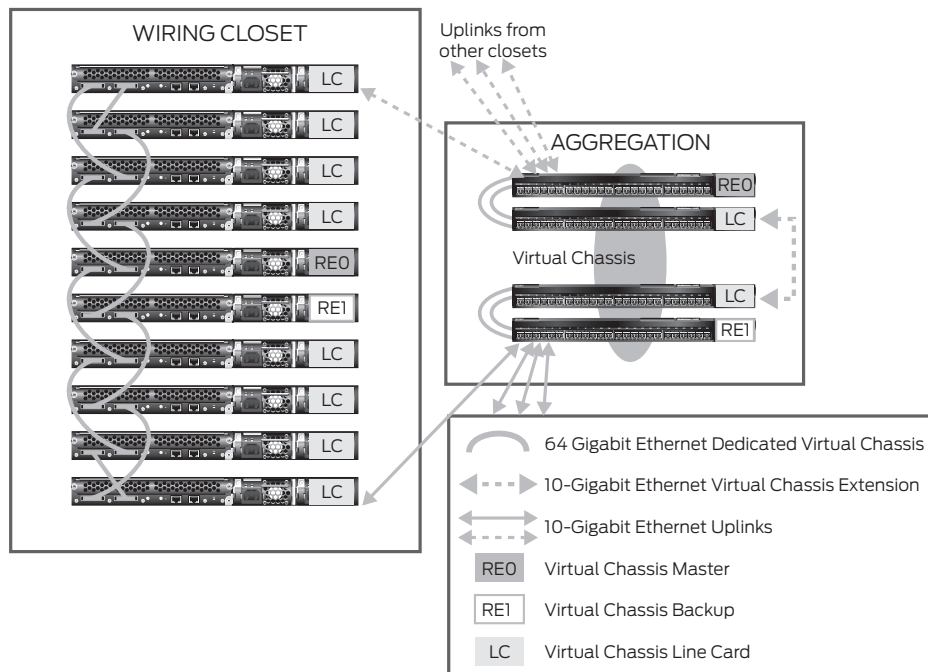


Figure 20: Campus aggregation, single Virtual Chassis configuration

Virtual Chassis Ring Deployment

The Virtual Chassis ring deployment is appropriate for metro Ethernet deployments and for campuses or retail operations. In this deployment, switches can be located in several different buildings.

Figure 21 illustrates an inter-building deployment. Figure 22 illustrates an intra-building deployment.

The Metro/Ring deployment provides benefits that include:

- One-switch LAN
- High availability
- Redundant power and fans
- Redundant switch fabric
- Sub-second convergence in case of a device or link failure
- Integrated access security
- Integrated quality of service (QoS) for voice, video, and data

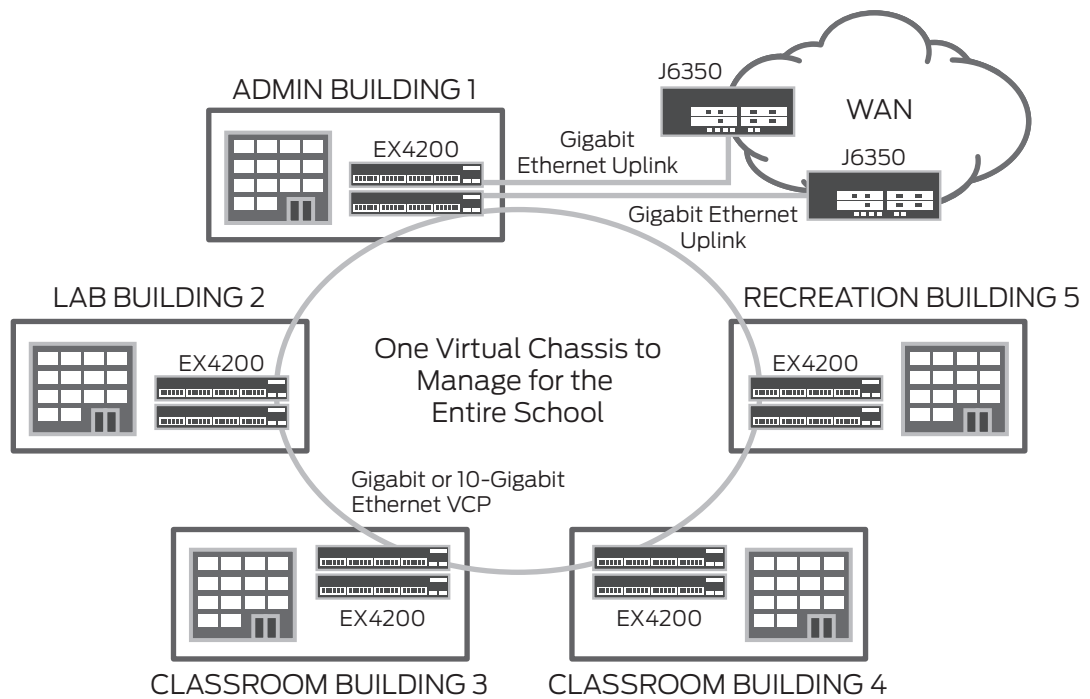


Figure 21: Virtual Chassis ring inter-building deployment example

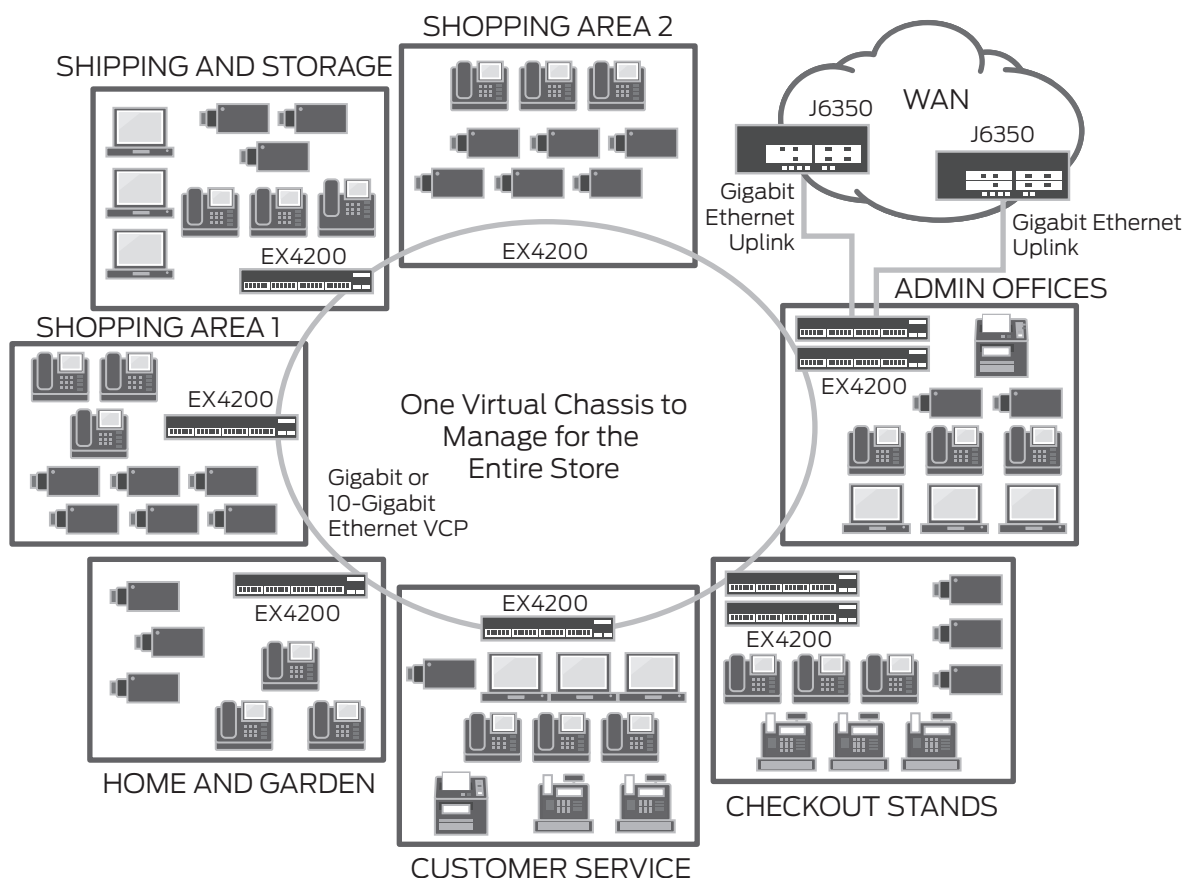


Figure 22: Virtual Chassis ring intra-building example

Best Practices

This section provides a summary of the best practices that we recommend you follow when deploying and operating the EX4200 Virtual Chassis technology.

- Use the VME interface as the management interface to configure Virtual Chassis technology options. For more information, please see the "Assigning an IP Address to a Virtual Chassis Configuration" section.
- When removing a switch from a Virtual Chassis configuration, immediately recycle its member ID so that the ID becomes the next lowest available unused ID. In this way, the replacement switch automatically is assigned that member ID and inherits the configuration of the original switch. For more information, please see the "Member ID Numbering" section.
- When installing a Virtual Chassis configuration, explicitly configure the mastership priority of the members that you want to function as the master and backup switches. For more information, please see the "Mastership Priority Setting" section.
- Specify the same mastership priority value for the master and backup switches in a Virtual Chassis configuration. For more information, please see the "Mastership Priority Setting" section.
- Configure the highest possible mastership priority value (255) for the master and backup switches. For more information, please see the "Mastership Priority Setting" section.
- When designing a Virtual Chassis configuration, consider a deployment in which ports are distributed across as many switches as possible to provide the highest resiliency and the smallest failure domain. For more information, please see the "How Many Members Should a Virtual Chassis Configuration Have?" section.
- Evenly space master and backup switches by member hop when possible. For more information, please see the "Location of Master and Backup Switches" section.
- Place the master and backup switches in separate locations when deploying an extended Virtual Chassis configuration. For more information, please see the "Location of Master and Backup Switches" section.
- For maximum resiliency, interconnect Virtual Chassis member devices in a ring topology. For more information, please see the "Cabling Options" section.

- When possible, place uplink modules in the Virtual Chassis configuration line card switches, and place uplinks in devices which are separated at equal distances by member hop. For more information, please see the “Using Uplinks” section.
- When changing configuration settings on the master switch, propagate changes to all other switches in the Virtual Chassis configuration. For more information, please see the “Synchronizing Virtual Chassis Members” section.

Conclusion

Juniper Networks Virtual Chassis technology allows the interconnection of up to 10 EX4200 switches to form a single, logical device that is managed as a single chassis. The Virtual Chassis technology provides the same level of availability, performance, and scale as a traditional modular chassis, with benefits that include smaller space and power requirements and a lower deployment cost. Additionally, the Virtual Chassis configuration provides flexibility in the physical deployment of devices across extended distances, enabling unique network designs to reduce deployment and operational costs. For additional information, please refer to the following documents:

- Branch Office Implementation Guide: www.juniper.net/us/en/local/pdf/implementation-guides/8010010-en.pdf
- Branch Reference Architecture: www.juniper.net/us/en/local/pdf/whitepapers/2000234-en.pdf
- Campus Reference Architecture: www.juniper.net/us/en/local/pdf/reference-architecture/8030005-en.pdf

Appendix: Glossary

Backup switch—The Virtual Chassis configuration member that takes over if the master switch stops operating.

Braided-ring—Method of cabling together switches in a Virtual Chassis configuration in which alternating devices connect to each other and the top two and bottom two devices, or pairs of device on each end in a horizontal configuration, connect to each other.

Daisy-chained ring—Method of cabling together switches in a Virtual Chassis configuration in which each device connects to the device immediately adjacent to it, and the top and bottom devices or the two end devices connect to each other.

Dedicated configuration—A Virtual Chassis configuration which consists of adjacent switches which are interconnected with Virtual Chassis port cables.

Extended Virtual Chassis configuration—Also known as “extended configuration,” a deployment which includes uplinks to interconnect individual Virtual Chassis members or dedicated Virtual Chassis configurations across distances of up to 50 km with redundant fiber links.

Forwarding path—Path for routing packets through switches in a Virtual Chassis configuration.

Graceful Routing Engine switchover (GRES)—Enables the backup switch to take over for the master switch with minimal interruption to network communications.

Junos OS for EX Series switches—Juniper Networks proprietary operating system on all EX Series switches and Juniper Networks routers.

J-Web—Graphical Web browser interface to the Junos OS Internet software on routing platforms.

Line card switch—Also known as “line card,” Virtual Chassis configuration member which is not assigned the master or backup role.

Link aggregation group (LAG)—A logical point-to-point link which is created by combining ports that belong to different member switches in a Virtual Chassis configuration.

Master switch—The Virtual Chassis configuration member that has full control over the Virtual Chassis configuration.

Mastership election process—Algorithm that the system runs to determine which member switches take the role of master switch, backup switch, and line cards.

Mastership priority—Value from 1 to 255 that the system uses when determining which switch functions in the role of master switch, backup switch, or line card.

Member switch—Also known as “member,” a switch that is part of a Virtual Chassis configuration.

Member ID—Number from 0 to 9 that distinguishes member switches from each other.

Packet Forwarding Engine (PFE)—Portion of the router that processes packets by forwarding them between input and output interfaces.

Pre-provisioning—A method for deterministically controlling the member ID and role which are assigned to a member switch by associating the switch to its serial number.

Routing Engine—Switch that calculates and maintains the forwarding table and provides it to the Packet Forwarding Engines (PFEs) in other member switches, and that receives and transmits routing information. The master switch is always the Routing Engine.

Uplink module—Optional add-on to a switch that allows the extension of a Virtual Chassis configuration across distances which are greater than the 5 meter limit imposed by a dedicated Virtual Chassis port cable.

Uplink port—Port on an uplink module which is used to interconnect switches in a Virtual Chassis configuration.

Uplinks—Enable the interconnection of remote switches, other Virtual Chassis members, and a variety of other devices to a Virtual Chassis configuration.

Virtual Chassis configuration—Implementation of Virtual Chassis technology across interconnected switches.

Virtual Chassis port—64 Gbps ports on EX4200 switches which are used exclusively to interconnect switches in a Virtual Chassis configuration. The ports are named VCP 0 and VCP 1.

Virtual Chassis port cable—Cable that interconnects switches through Virtual Chassis ports.

Virtual Chassis technology—Feature in Juniper Networks EX4200 line of switches that allows you to interconnect switches and operate them as a unified, single, high bandwidth device.

Virtual Management Ethernet (VME) interface—Enables the remote management of the Virtual Chassis configuration by connecting to the out-of-band management port of any member switch through a single IP address.

About Juniper Network

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.

Corporate and Sales Headquarters

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters

Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters

Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2010 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

8010018-006-EN Feb 2010

 Printed on recycled paper