

Understanding Split and Merge in a Virtual Chassis Configuration

In a Virtual Chassis configuration, two or more Juniper Networks EX4200 Ethernet Switches are connected together to form a unit that is managed as a single chassis. If there is a disruption to the Virtual Chassis configuration due to member switches failing or being removed from the configuration, the Virtual Chassis configuration splits into two separate Virtual Chassis. This situation could cause disruptions in the network if the two separate configurations share common resources, such as global IP addresses. The split and merge feature provides a method to prevent the separate Virtual Chassis configurations from adversely affecting the network and also allows the two parts to merge back into a single Virtual Chassis configuration.



NOTE: If a Virtual Chassis configuration splits into separate parts, we recommend that you resolve the problem that caused the Virtual Chassis configuration to split as soon as possible.

You can also use this feature to merge two active but separate Virtual Chassis that have not previously been part of the same configuration into one Virtual Chassis configuration.



NOTE: The split and merge feature is enabled by default on EX4200 switches. You can disable the split and merge feature by using the `set virtual-chassis no-split-detection` command.

This topic describes:

- What Happens When a Virtual Chassis Configuration Splits on page 1
- Merging Virtual Chassis Configurations on page 2

What Happens When a Virtual Chassis Configuration Splits

When a Virtual Chassis configuration splits into two separate Virtual Chassis configurations, the individual member switches detect this topology change and run the master election algorithm to select a new master for each of the two Virtual Chassis configurations. The new masters then determine whether their Virtual Chassis configuration remains active. One of the configurations remains active based on the following:

- It contains both the stable master and the stable backup (that is, the master and backup from the original Virtual Chassis configuration before the split).
- It contains the stable master and the configuration is greater than half the Virtual Chassis size.
- It contains the stable backup and is at least half the Virtual Chassis size.

Due to the rules given in the second and third list items, if the Virtual Chassis configuration splits into two equal parts and the stable master and stable backup are in different parts, then the part that contains the stable backup will become active.



NOTE: The number of members in the Virtual Chassis configuration includes all member switches connected to date minus the number whose Virtual Chassis member IDs have been recycled. Therefore, the size of the Virtual Chassis configuration increases when a new member switch is detected and decreases when a member switch's ID is recycled (that is, made available for reassignment).

These rules ensure that only one of the two separate Virtual Chassis configurations created by the split remains active. The member switches in the inactive Virtual Chassis configuration remain in a linecard role. For the inactive members to become active again, one of the following things must happen:

- The problem that caused the original Virtual Chassis configuration to split is resolved, allowing the two Virtual Chassis configurations to merge.
- You load the factory default configuration on the inactive members, which causes the inactive members to function as standalone switches or become part of a different Virtual Chassis configuration.



NOTE: When you remove a member switch from a Virtual Chassis configuration, you should recycle the member ID using the `request virtual-chassis recycle` command.

Merging Virtual Chassis Configurations

There are two scenarios in which separate Virtual Chassis merge:

- A Virtual Chassis configuration that had split into two is now merging back into a single configuration because the problem that had caused it to split has been resolved.
- You want to merge two Virtual Chassis that had not previously been configured together.

Every Virtual Chassis configuration has a unique ID that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a Virtual Chassis ID using the `set virtual-chassis id` command. A Virtual Chassis ID that you assign takes precedence over automatically assigned Virtual Chassis IDs.

When you reconnect the separate Virtual Chassis configurations or connect them for the first time, the members determine whether or not the separate Virtual Chassis configurations can merge. The members use the following rules to determine whether a merge is possible:

- If the Virtual Chassis configurations have the same Virtual Chassis ID, then the configurations can merge. If the two Virtual Chassis were formed as the result of a split, they will have the same Virtual Chassis ID.

- If the Virtual Chassis IDs are different, then the two configurations can merge only if both are active (inactive configurations cannot merge, ensuring that members removed from one Virtual Chassis configuration do not become members of another Virtual Chassis configuration). If the configurations to merge are both active and one of them has a user-configured Virtual Chassis ID, this ID becomes the ID of the merged Virtual Chassis. If neither Virtual Chassis has a user-configured Virtual Chassis ID, then the Virtual Chassis ID of the configuration with the highest mastership priority becomes the ID of the merged Virtual Chassis. The resulting merged Virtual Chassis configuration will be active.

When you connect two Virtual Chassis configurations, the following events occur:

1. Connecting the two split Virtual Chassis configurations triggers the shortest-path-first (SPF) algorithm. The SPF algorithm computes the network topology and then triggers the master election algorithm. The master election algorithm waits for the members to synchronize the topology information before running.
2. The master election algorithm merges the Virtual Chassis IDs of all the members.
3. Each member runs the master election algorithm to select a master and a backup from among all members with the same Virtual Chassis IDs. For more information, see *Understanding How the Master in a Virtual Chassis Configuration Is Elected*.
4. The master determines whether the Virtual Chassis configuration is active or inactive. (See “What Happens When a Virtual Chassis Configuration Splits” on page 1.)
5. If the Virtual Chassis configuration is active, the master assigns roles to all members. If the Virtual Chassis configuration is inactive, the master assigns all members the role of linecard.
6. When the other members receive their role from the master, they change their role to backup or linecard. They also use the active or inactive state information sent by the master to set their own state to active or inactive and to construct the Virtual Chassis member list from the information sent by the master.
7. If the Virtual Chassis state is active, the master waits for messages from the members indicating that they have changed their roles to the assigned roles, and then the master changes its own role to master.



NOTE: When you merge two Virtual Chassis that had not previously been part of the same Virtual Chassis configuration, any configuration settings (such as the settings for Telnet/FTP services, GRES, fast failover, VLANs, and so on) that exist on the new master will become the configuration settings for all members of the new Virtual Chassis, overwriting any other configuration settings.

Related Topics

- Understanding Virtual Chassis Configuration
- Example: Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge

- Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure)
- Disabling Split and Merge in a Virtual Chassis Configuration (CLI Procedure)

Published: 2009-07-29