High Availability Features for EX-series Switches Overview

High availability refers to the hardware and software components that provide redundancy and reliability for packet-based communications. This topic covers the following high availability features of EX-series switches:

VRRP

For Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and logical interfaces on EX-series switches, you can configure the Virtual Router Redundancy Protocol (VRRP). The switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup EX-series switch can take over a failed default switch within few seconds. This is done with minimum VRRP traffic and without any interaction with the hosts.

NOTE: The VRRP master and backup routing platforms should not be confused with the master and backup member switches of a Virtual Chassis configuration. The master and backup members of a Virtual Chassis configuration compose a single host. In a VRRP topology, one host operates as a master routing platform and another host operates as a backup routing platform, as shown in Figure 2 on page 2.

Switches running VRRP dynamically elect master and backup routing platforms. You can also force assignment of master and backup routing platforms using priorities from 1 through 255, with 255 being the highest priority. In VRRP operation, the default master routing platform sends advertisements to backup routing platforms at regular intervals. The default interval is 1 second. If a backup routing platform does not receive an advertisement for a set period, the backup routing platform with the next highest priority takes over as master and begins forwarding packets.

Figure 1 on page 2 illustrates a basic VRRP topology with EX-series switches. In this example, Switches A, B, and C are running VRRP and together they make up a virtual routing platform. The IP address of this virtual routing platform is 10.10.0.1 (the same address as the physical interface of Switch A).

Figure 1: Basic VRRP on EX-series Switches

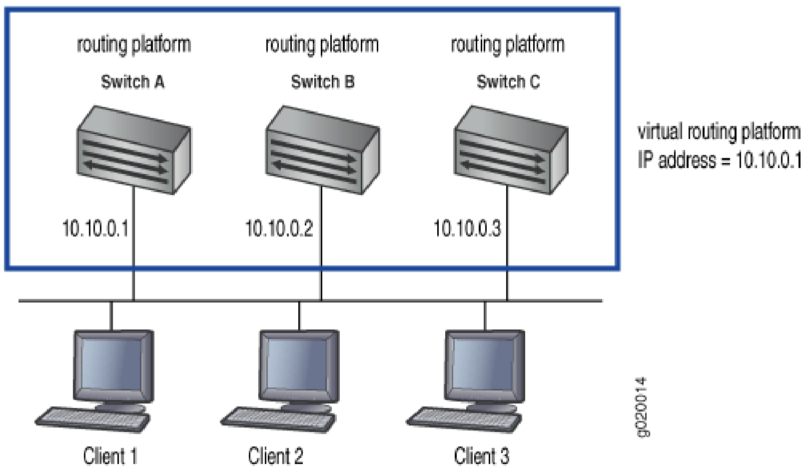

Figure 2 on page 2 illustrates a basic VRRP topology using Virtual Chassis configurations. Switch A, Switch B, and Switch C are each composed of multiple interconnected EX 4200 switches. Each Virtual Chassis configuration operates as a single switch, which is running VRRP, and together they make up a virtual routing platform. The IP address of this virtual routing platform is 10.10.0.1 (the same address as the physical interface of Switch A).

Figure 2: VRRP on EX 4200 Virtual Chassis Switches

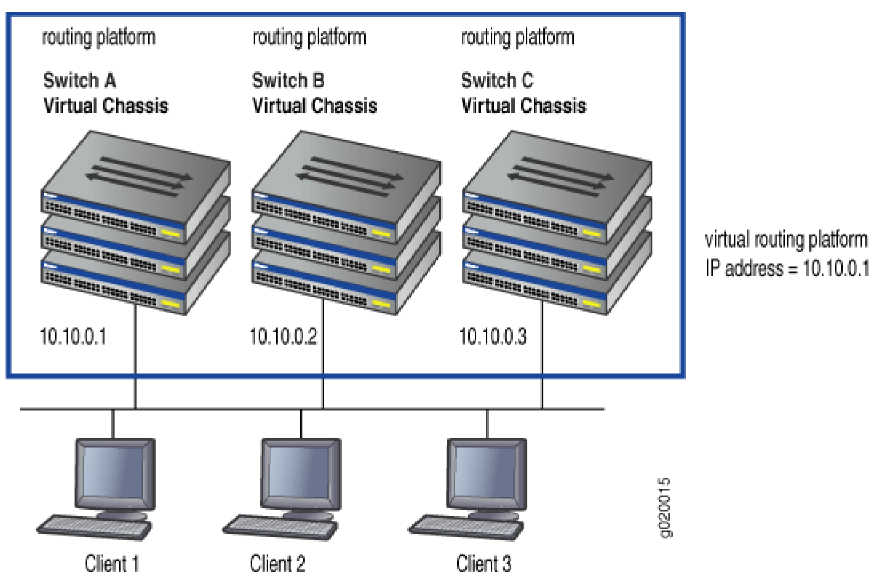

Because the virtual routing platform uses the IP address of the physical interface of Switch A, Switch A is the master VRRP routing platform, while switches B and C function as backup VRRP routing platforms. Clients 1 through 3 are configured with the default gateway IP address of 10.10.0.1. As the master router, Switch A forwards packets sent to its IP address. If the master virtual routing platform fails, the switch configured with the higher priority becomes the master virtual routing platform and

provides uninterrupted service for the LAN hosts. When Switch A recovers, it becomes the master virtual routing platform again.

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

Graceful Protocol Restart

With standard implementations of routing protocols, any service interruption requires an affected switch to recalculate adjacencies with neighboring switches, restore routing table entries, and update other protocol-specific information. An unprotected restart of a switch can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. Graceful protocol restart allows a restarting switch and its neighbors to continue forwarding packets without disrupting network performance. Because neighboring switches assist in the restart (these neighbors are called helper switches), the restarting switch can quickly resume full operation without recalculating algorithms from scratch.

On EX-series switches, graceful protocol restart can be applied to aggregate and static routes and for routing protocols (BGP, IS-IS, OSPF and RIP).

Graceful protocol restart works similarly for the different routing protocols. The main benefits of graceful protocol restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful protocol restart thus allows a switch to pass through intermediate convergence states that are hidden from the rest of the network. Most graceful restart implementations define two types of switches—the restarting switch and the helper switch. The restarting switch requires rapid restoration of forwarding state information so it can resume the forwarding of network traffic. The helper switch assists the restarting switch in this process. Individual graceful restart configuration statements typically apply to either the restarting switch or the helper switch.

EX 4200 Redundant Routing Engines

Two to ten EX 4200 switches can be interconnected to create a Virtual Chassis configuration that operates as a single network entity. Every Virtual Chassis configuration with two or more members has a master and a backup. The master acts as the master Routing Engine and the backup acts as the backup Routing Engine. The Routing Engine provides the following functionality:

- Runs various routing protocols
- Provides the forwarding table to the Packet Forwarding Engines (PFEs) in all the member switches of the Virtual Chassis configuration
- Runs other management and control processes for the entire Virtual Chassis configuration

The master Routing Engine, which is in the master of the Virtual Chassis configuration, runs JUNOS software in the master role. It receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding Engine components of the member switches, and has full control over the Virtual Chassis configuration.

The backup Routing Engine, which is in the backup of the Virtual Chassis configuration, runs JUNOS software in a backup role. It stays in sync with the master Routing Engine in terms of protocol states, forwarding tables, and so forth. If the master becomes unavailable, the backup Routing Engine takes over the functions that the master Routing Engine performs.

EX 4200 Graceful Routing Engine Switchover

You can configure graceful Routing Engine switchover (GRES) in a Virtual Chassis configuration, allowing the configuration to switch from the master Routing Engine in the master to the backup Routing Engine in the backup with minimal interruption to network communications. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and forwarding state. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.

When the backup Routing Engine assumes mastership in a redundant failover configuration (when graceful Routing Engine switchover is not enabled), the Packet Forwarding Engines initialize their state to boot up state before they connect to the new master Routing Engine. In contrast, in a graceful switchover configuration, the Packet Forwarding Engines do not reinitialize their state, but instead resynchronize their state with the new master Routing Engine. The interruption to the traffic is minimal.

Graceful Routing Engine switchover on EX 4200 switches supports software features in JUNOS Release 9.2 or later for EX-series switches.

EX 4200 Virtual Chassis Software Upgrade and Failover Features

EX 4200 switches provide these features for increased resiliency in Virtual Chassis configurations:

- Virtual Chassis atomic software upgrade—When you upgrade software in a Virtual Chassis configuration, the upgrade will either succeed or fail on all member switches, preventing the situation in which only some Virtual Chassis member switches are upgraded.
- Virtual Chassis fast failover—A hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link failure.
- Virtual Chassis split and merge—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.

Link Aggregation

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a *link aggregation group (LAG)* or *bundle*. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation

provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.

You can select up to eight Ethernet interfaces and include them within a link aggregation group. In an EX 4200 Virtual Chassis configuration composed of multiple members, the interfaces that compose a LAG can be on different members of the Virtual Chassis. See Understanding Virtual Chassis Configurations and Link Aggregation.

Additional High Availability Features of EX-series Switches

To ensure continuous operation, all EX-series switches use field-replaceable power supply units, fan trays, and uplink modules. EX 4200 switches include options for external power-supply redundancy.

The EX 3200 switches support a single field-replaceable power supply unit, a field-replaceable fan tray, and a field-replaceable uplink module.

The EX 4200 switches supports connection of virtual chassis members using two dedicated Virtual Chassis ports (VCPs) on the rear panel or SFP uplink module ports. The EX 4200 switches also support two internal load-sharing redundant hot-swappable power supplies, field-replaceable fan trays with redundant blowers, and field-replaceable uplink modules that provide SFP or XFP ports.

Notification of hardware issues is provided through system log messages and alarms.

Related Topics

- For more information on high availability features, see the *JUNOS Software High Availability Configuration Guide* at http://www.juniper.net/techpubs/software/junos/junos90/index.html.
- Virtual Chassis Overview
- Understanding Virtual Chassis Components
- Understanding Virtual Chassis Configurations and Link Aggregation