Understanding Generic Routing Encapsulation

www.juniper.net/techpubs/en_US/junos12.1/topics/concept/gre-tunnel-services.html

Tunneling provides a private, secure path for transporting packets through an otherwise public network by encapsulating packets inside a transport protocol known as an *IP encapsulation protocol*. Generic routing encapsulation (GRE) is an IP encapsulation protocol that is used to transport packets over a network. Information is sent from one network to the other through a GRE tunnel.

Overview of GRE

GRE encapsulates packets into IP packets and redirects them to an intermediate host, where they are de-encapsulated and routed to their final destination. Because the route to the intermediate host appears to the inner datagrams as one hop, Juniper Networks EX Series Ethernet switches can operate as if they have a virtual point-to-point connection with each other. GRE tunnels allow routing protocols like RIP and OSPF to forward data packets from one switch to another switch across the Internet. In addition, GRE tunnels can encapsulate multicast data streams for transmission over the Internet.

GRE is described in RFC 2784 (obsoletes earlier RFCs 1701 and 1702). The switches support RFC 2784, but not completely. (For a list of limitations, see Configuration Limitations for GRE on EX Series Switches.)

As a *tunnel source router*, the switch encapsulates a payload packet for transport through the tunnel to a destination network. The payload packet is first encapsulated in a GRE packet, and the resulting GRE packet is encapsulated in a delivery protocol. The switch performing the role of a *tunnel remote router* extracts the tunneled packet and forwards the packet to the destination network.

GRE Tunneling

Data is routed by the system to the GRE endpoint over routes established in the route table. (These routes can be statically configured or dynamically learned by routing protocols such as RIP or OSPF.) When a data packet is received by the GRE endpoint, it is de-encapsulated and routed again by means of the endpoint configuration to the destination address of the tunnel. In this way, each data packet traveling over the GRE tunnel gets routed through the system twice.

Because GRE tunnels are *stateless*, the endpoint of the tunnel contains no information about the state or availability of the remote tunnel endpoint. Therefore, the switch operating as a tunnel source router cannot change the state of the GRE tunnel interface to down if the remote endpoint is unreachable.

For details of GRE tunneling, see:

Encapsulation and De-Encapsulation on the Switch

Encapsulation—A switch operating as a tunnel source router encapsulates and forwards GRE packets as follows:

1. When a switch receives a data packet (payload) to be tunneled, it sends the packet to the tunnel interface.
2. The tunnel interface encapsulates the data in a GRE packet.
3. The system encapsulates the GRE packet in an IP packet.
4. The IP packet is forwarded based on its destination address and routing table.

De-encapsulation—A switch operating as a tunnel remote router handles GRE packets as follows:

1. When the destination switch receives the IP packet from the tunnel interface, the switch checks the destination address.

2. The IP header is removed, and the packet is submitted to the GRE protocol.
3. The GRE protocol strips off the GRE header and submits the payload packet for forwarding.

Number of Source and Destination Tunnels Allowed on a Switch

Depending on your network, you can configure up to approximately 500 GRE tunnels to operate between switches transmitting IPv4 or IPv6 payload packets over GRE. If a passenger protocol in addition to IPv4 and IPv6 is used, you can configure up to approximately 333 GRE tunnels between the switches.

A switch can have a maximum of 20 tunnel source IP addresses configured, and each tunnel source IP can be configured with up to 20 destination IP addresses on a second switch. As a result, the two connected switches can have a maximum of 400 GRE tunnels. If the first switch is also connected to a third switch, the possible maximum number of tunnels can reach 500.

Class of Service on GRE Tunnels

When a network experiences congestion and delay, some packets might be dropped. Junos OS class of service (CoS) divides traffic into classes to which you can apply different levels of throughput and packet loss when congestion occurs and thereby set rules for packet loss. For CoS details, see Junos OS CoS for EX Series Switches Overview.

The following CoS components are available on a switch operating as a GRE tunnel source router or GRE tunnel remote router:

- At the GRE tunnel source—On a switch operating as a tunnel source router, you can apply CoS classifiers on an ingress port or on a GRE port, with the following results on CoS component support on tunneled packets:
  - Schedulers only—Based on the CoS classification on the *ingress port*, you can apply CoS schedulers on a GRE port of the switch to define output queues and control the transmission of packets through the tunnel after GRE encapsulation. However, you cannot apply CoS rewrite rules to these packets.
  - Schedulers and rewrite rules—Based on the CoS classification on the *GRE port*, you can apply both schedulers and rewrite rules to the encapsulated packets transmitted through the tunnel.
- At the GRE tunnel endpoint—When the switch is a tunnel remote router, you can apply CoS classifiers on the GRE port and schedulers and rewrite rules on the egress port to control the transmission of a de-encapsulated GRE packet out the egress port.

Firewall Filters on GRE Tunnels

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a switch from a source address to a destination address. For details, see Firewall Filters for EX Series Switches Overview.

The effectiveness of a firewall filter on a packet sent through a GRE tunnel depends on whether a switch is at the source or endpoint of the tunnel:

- GRE tunnel source—On a switch operating as a tunnel source router, you can configure a firewall filter on a egress port to control tunnel payload packets and GRE-encapsulated IP packets at the tunnel source.
- GRE tunnel endpoint—When the switch is a tunnel remote router, a firewall filter can control a GRE packet only after the packet has been de-encapsulated.

Configuration Limitations for GRE on EX Series Switches

Some GRE tunneling features are not currently available on EX Series switches. Be aware of the following limitations when you are configuring GRE on a switch:

- Unsupported features—GRE on the switches *does not support* the following features:
  - Virtual routing over GRE
  - Bidirectional Forwarding Detection (BFD) protocol over GRE distributed mode
  - MPLS over GRE tunnels
  - GRE keepalives
  - GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets
  - BGP dynamic tunnels
- Platform limitation—Encapsulation over GRE tunnels is supported on EX3200, EX4200, and EX8200 switches. However, de-encapsulation is supported only on EX3200 and EX4200 switches.
- IPv6 limitation—Both IPv4 and IPv6 packets are accepted as payload packets for GRE encapsulation, but only IPv4 is supported as the GRE delivery protocol (the protocol through the tunnel).
- OSPF limitation—Enabling OSPF on a GRE interface creates two equal-cost routes to the destination: one through the Ethernet network or uplink interface (for example, ge-*fpc*/*pic*/*port*) and the other through the tunnel interface (gr-*fpc*/*pic*/*port*). If data is routed through the tunnel interface, the interface might go down. To keep the interface operational, we recommend that you use a static route, disable OSPF on the tunnel interface, or configure the peer not to advertise the tunnel destination over the tunnel interface.

Related Documentation

- Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)

Published: 2012-06-06