Understanding Q-in-Q Tunneling on EX-series Switches

Q-in-Q tunneling allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag. The JUNOS software implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

This topic describes:

How Q-in-Q Tunneling Works

In Q-in-Q tunneling, as a packet travels from a customer VLAN (C-VLAN) to a service provider's VLAN, a customer-specific 802.1Q tag is added to packets. This additional tag is used to segregate traffic into service–provider–defined service VLANs. The original customer 802.1Q tag of the packet remains and is transmitted transparently, passing through the service provider's network. The S-VLAN tag is added on egress for incoming packets, optionally including untagged packes. As the packet leaves the S-VLAN in the downstream direction, the extra 802.1Q tag is removed, leaving the original customer tag on the packet.

When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider network and access interfaces are assumed to be customer facing. An access interface can receive both tagged and untagged frames in this case.

A trunk interface can be a member of multiple S-VLANs. You can map one C-VLAN to one S-VLAN (1:1) or multiple C-VLANs to one S-VLAN (N:1). You can also double-tag packets for an additional layer of segregating or bundling of C-VLANs. C-VLAN and S-VLAN tags are unique, so you can have both C-VLAN 101 and S-VLAN 101, for example. You can limit the set of accepted customer tags to a list of ranges or discrete values. Class-of-service (CoS) values of C-VLANs are unchanged in the downstream direction. You may, optionally, copy ingress priority and CoS settings to the S-VLAN. Using private VLANs, you can isolate users to prevent forwarding traffic between user interfaces even if the interfaces are on the same VLAN.

Disabling MAC Address Learning

In a Q-in-Q deployment, customer packets from downstream interfaces are transported without any changes to the source and destination MAC addresses. The service provider switches might learn a large number of MAC addresses, and that abundance of learned MAC addresses might slow performance.. You can disable MAC address learning at both the interface level and the VLAN level. Disabling MAC address learning on an interface disables learning for all the VLANs of which that

interface is a member. When you disable MAC address learning on a VLAN, MAC addresses that have already been learned are flushed.

If you disable MAC address learning on an interface or a VLAN, you cannot include MAC move limiting or 802.1X authentication in that same VLAN configuration.

When a routed VLAN interface (RVI) is associated with either an interface or a VLAN on which MAC address learning is disabled, the Layer 3 routes resolved on that VLAN or that interface are not resolved with the Layer 2 component. This results in routed packets flooding all the interfaces associated with the VLAN.

Limitations for Q-in-Q Tunneling

You cannot add the C-VLAN tag on egress for incoming untagged packets or remove the C-VLAN tag in the downstream direction. Q-in-Q does not support IGMP snooping or most access port security features. There is no per-VLAN (customer) policing or per-VLAN (outgoing) shaping and limiting with Q-in-Q.

Related Topics

- Understanding Bridging and VLANs on EX-series Switches
- Example: Setting Up Q-in-Q Tunneling on EX-series Switches
- Configuring Q-in-Q Tunneling (CLI Procedure)