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 Juniper Networks 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 the packet. This additional tag is used to segregate traffic into service-provider-defined service VLANs (S-VLANs). The original customer 802.1Q tag of the packet remains and is transmitted transparently, passing through the service provider's network. As the packet leaves the S-VLAN in the downstream direction, the extra 802.1Q tag is removed.

When Q-in-Q tunneling is enabled on Juniper Networks EX Series Ethernet Switches, 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.

An 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). Packets are double-tagged for an additional layer of segregating or bundling of C-VLANs. C-VLAN and S-VLAN tags are unique; so you can have both a C-VLAN 101 and an S-VLAN 101, for example. You can limit the set of accepted customer tags to a range of tags or to 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 the forwarding of traffic between user interfaces even if the interfaces are on the same VLAN.

You can use the native option to specify an S-VLAN for untagged and priority tagged packets when using many-to-one bundling and mapping a specific interface approaches to map C-VLANs to S-VLANs. Otherwise the packets are discarded. The native option is not available for all-in-one bundling because there is no need to

specify untagged and priority tagged packets when all packets are mapped to the C-VLAN. See the Mapping C-VLANs to S-VLANs section of this document for information on the methods of mapping C-VLANs to S-VLANs.

Firewall filters allow you to map an interface to a VLAN based on a policy. Using firewall filters to map an interface to a VLAN is useful when you want a subset of traffic from a port to be mapped to a selected VLAN instead of the designated VLAN. To configure a firewall filter to map an interface to a VLAN, the vlan option has to be configured as part of the firewall filter and the mapping policy option must be specified in the interface configuration for each logical interface using the filter.

Disabling MAC Address Learning

In a Q-in-Q deployment, customer packets from downstream interfaces are transported without any changes to source and destination MAC addresses. 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.

Mapping C-VLANs to S-VLANs

There are three ways to map C-VLANs to an S-VLAN:

- All-in-one bundling—Use the dot1q-tunneling option to map without specifying customer VLANs. All packets from all access interfaces are mapped to the S-VLAN.
- Many-to-one bundling—Use the customer-vlans option to specify which C-VLANs are mapped to the S-VLAN.
- Mapping a specific interface—Use the mapping option to indicate a specific S-VLAN for a given C-VLAN. The specified C-VLAN applies to only one VLAN and not all access interfaces as in the cases of all-in-one and many-to-one bundling.

If you configure multiple methods, the switch gives priority to mapping a specific interface, then to many-to-one bundling, and last to all-in-one bundling. However, you cannot have overlapping rules for the same C-VLAN under a given approach.

All-in-One Bundling

All-in-one bundling maps all packets from all access interfaces to the S-VLAN. All-in-one bundling is configured using the dot1q-tunneling option without specifying customer VLANs.

When all-in-one bundling is used, all packets leaving the C-VLAN, including untagged and priority tagged packets, enter the S-VLAN.

Many-to-One Bundling

Many-to-one bundling is used to specify which C-VLANs are mapped to an S-VLAN. Many-to-one bundling is configured using the customer-vlans option.

Many-to-one bundling is used when you want a subset of the C-VLANs on the access switch to be part of the S-VLAN. When using many-to-one bundling, untagged and priority tagged packets can be mapped to the S-VLAN when the native option is specified along with the customer-vlans option.

Mapping a Specific Interface

Use the mapping a specific interface approach when you want to assign an S-VLAN to a specific C-VLAN on an interface. The mapping a specific interface configuration only applies to the configured interface, not to all access interfaces as in the cases of the all-in-one bundling and many-to-one bundling approaches. The mapping a specific interface approach is configured using the mapping option to indicate a specific S-VLAN for a given C-VLAN.

It might be useful to have S-VLANs that provide service to multiple customers. Each customer will typically have its own S-VLAN plus access to one or more S-VLANs that are used by multiple customers. A specific tag on the customer side is mapped to an S-VLAN. Typically, this functionality is used to keep data from different customers separate or to provide individualized treatment of the packets on a certain interface.

Routed VLAN Interfaces on Q-in-Q VLANs

Routed VLAN interfaces (RVIs) are supported on Q-in-Q VLANs.

Packets arriving on an RVI that is using Q-in-Q VLANs will get routed regardless of whether the packet is single or double tagged. The outgoing routed packets contain an S-VLAN tag only when exiting a trunk interface; the packets exit the interface untagged when exiting an access interface.

Limitations for Q-in-Q Tunneling

Q-in-Q tunneling does not support most access port security features. There is no per-VLAN (customer) policing or per-VLAN (outgoing) shaping and limiting with Q-in-Q tunneling unless you configure these security features using firewall filters.

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)

Published: 2009-07-22