



Cisco ACI QoS

This chapter contains the following sections:

- [QoS for L3Outs, on page 1](#)
- [Preserving 802.1P Class of Service Settings, on page 1](#)
- [Preserving QoS Priority Settings in a Multipod Fabric, on page 2](#)
- [Translating QoS Ingress Markings to Egress Markings, on page 3](#)

QoS for L3Outs

To configure QoS policies for an L3Out, use the following guidelines:

- To configure the QoS policy to be enforced on the border leaf where the L3Out is located, the VRF must be in egress mode (Policy Control Enforcement Direction must be "Egress").
- To enable the QoS policy to be enforced, the VRF Policy Control Enforcement Preference must be "Enforced."
- When configuring the contract governing communication between the L3Out and other EPGs, include the QoS class or Target DSCP in the contract or subject.



Note Only configure a QoS class or target DSCP in the contract, not in the external EPG (`l3extInstP`).

Preserving 802.1P Class of Service Settings

APIC enables preserving 802.1P class of service (CoS) settings within the fabric. Enable the fabric global QoS policy `dot1p-preserve` option to guarantee that the CoS value in packets which enter and transit the ACI fabric is preserved.

802.1P CoS preservation is supported in single pod and multipod topologies.

In multipod topologies, CoS Preservation can be used where you want to preserve the QoS priority settings of 802.1P traffic entering POD 1 and egressing out of POD 2, but you are not concerned with preserving the CoS/DSCP settings in interpod network (IPN) traffic between the pods. To preserve CoS/DSCP settings when

multipod traffic is transitting an IPN, use a DSCP policy. For more information, see [Preserving QoS Priority Settings in a Multipod Fabric, on page 2](#).

Observe the following 802.1P CoS preservation guidelines and limitations:

- The current release can only preserve the 802.1P value within a VLAN header. The DEI bit is not preserved.
- For VXLAN encapsulated packets, the current release will not preserve the 802.1P CoS value contained in the outer header.
- 802.1P is not preserved when the following configuration options are enabled:
 - Multipod QoS (using a DSCP policy) is enabled.
 - Contracts are configured that include QoS.
 - Dynamic packet prioritization is enabled.
 - The outgoing interface is on a FEX.
- Preserving QoS CoS priority settings is not supported when traffic is flowing from an EPG with isolation enforced to an EPG without isolation enforced.
- A DSCP QoS policy is configured on a VLAN EPG and the packet has an IP header. DSCP marking can be set at the filter level on the following with the precedence order from the innermost to the outermost:
 - Contract
 - Subject
 - In Term
 - Out Term



Note When specifying vzAny for a contract, external EPG DSCP values are not honored because vzAny is a collection of all EPGs in a VRF, and EPG specific configuration cannot be applied. If EPG specific target DSCP values are required, then the external EPG should not use vzAny.


Preserving QoS Priority Settings in a Multipod Fabric

This topic describes how to guarantee QoS priority settings in a multipod topology, where devices in the interpod network are not under APIC management, and may modify 802.1P settings in traffic transitting their network.



Note You can alternatively use CoS Preservation where you want to preserve the QoS priority settings of 802.1P traffic entering POD 1 and egressing out of POD 2, but you are not concerned with preserving the CoS/DSCP settings in interpod network (IPN) traffic between the pods. For more information, see [Preserving 802.1P Class of Service Settings, on page 1](#).

Figure 1: Multipod Topology



As illustrated in this figure, traffic between pods in a multipod topology passes through an IPN, which may not be under APIC management. When an 802.1P frame is sent from a spine or leaf switch in POD 1, the devices in the IPN may not preserve the CoS setting in 802.1P frames. In this situation, when the frame reaches a POD 2 spine or leaf switch, it has the CoS level assigned by the IPN device, instead of the level assigned at the source in POD 1. Use a DSCP policy to ensure that the QoS priority levels are preserved in this case.

Configure a DSCP policy to preserve the QoS priority settings in a multipod topology, where there is a need to do deterministic mapping from CoS to DSCP levels for different traffic types, and you want to prevent the devices in the IPN from changing the configured levels. With a DSCP policy enabled, APIC converts the CoS level to a DSCP level, according to the mapping you configure. When a frame is sent from POD 1 (with the PCP level mapped to a DSCP level), when it reaches POD 2, the mapped DSCP level is then mapped back to the original PCP CoS level.

Translating QoS Ingress Markings to Egress Markings

APIC enables translating the 802.1P CoS field (Class of Service) based on the ingress DSCP value. 802.1P CoS translation is supported only if DSCP is present in the IP packet and dot1P is present in the Ethernet frames.

This functionality enables the ACI Fabric to classify the traffic for devices that classify the traffic based only on the CoS value. It allows mapping the dot1P CoS value based on the ingress dot1P value. It is mainly applicable for Layer 2 packets, which do not have an IP header.

Observe the following 802.1P CoS translation guidelines and limitations:

- Enable the fabric global QoS policy `dot1p-preserve` option.

- 802.1P CoS translation is not supported on external L3 interfaces.
- 802.1P CoS translation is supported only if the egress frame is 802.1Q encapsulated.

802.1P CoS translation is not supported when the following configuration options are enabled:

- Contracts are configured that include QoS.
- The outgoing interface is on a FEX.
- Multipod QoS using a DSCP policy is enabled.
- Dynamic packet prioritization is enabled.
- If an EPG is configured with intra-EPG endpoint isolation enforced.
- If an EPG is configured with allow-microsegmentation enabled.