



Configuring BGP EVPN VXLAN QoS

BGP EVPN VXLAN QoS enables you to provide Quality of Service (QoS) capabilities to traffic that is EVPN VXLAN-encapsulated.

By configuring QoS you can provide preferential treatment to specific types of traffic at the expense of other traffic types. Without QoS, the device offers best-effort service for each packet, regardless of the packet contents or size; the device sends the packets without any assurance of reliability, delay bounds, or throughput.

QoS provides the following features:

- Low latency
- Bandwidth guarantee
- Buffering capabilities and dropping disciplines
- Traffic policing
- Enables the changing of the attribute of the frame or packet header
- Relative services

For more information about QoS and its configuration, see *Quality of Service Configuration Guide, Cisco IOS XE* for the device.

- [Quality of Service Components, on page 1](#)
- [How to Configure BGP EVPN VXLAN QoS, on page 2](#)
- [BGP EVPN VXLAN QoS Modes, on page 3](#)
- [BGP EVPN VXLAN Traffic without QoS Policy, on page 4](#)
- [BGP EVPN VXLAN QoS Marking with an Input Policy, on page 5](#)
- [BGP EVPN VXLAN QoS Marking with Input and Output Policies, on page 6](#)

Quality of Service Components

QoS consists of the following key components:

QoS Component	Description
Classification	Classification is the process of distinguishing one type of traffic from another based upon access control lists (ACLs), Differentiated Services Code Point (DSCP), Class of Service (CoS), and other factors.

QoS Component	Description
Marking	<p>Marking is used on traffic to convey specific information to a downstream device in the network, or to carry information from one interface in a device to another. When traffic is marked, QoS operations on that traffic can be applied. This can be accomplished directly using the set command or through a table map, which takes input values and translates them directly to values on output.</p> <p>Through marking, you can set: a Layer 2 COS value in a frame, Layer 3 DSCP value in a packet, and Layer 3 ECN value in a packet.</p>
Shaping and Policing	<p>Shaping is the process of imposing a maximum rate of traffic, while regulating the traffic rate in such a way that downstream devices are not subjected to congestion. Shaping in the most common form is used to limit the traffic sent from a physical or logical interface. Policing is used to impose a maximum rate on a traffic class. If the rate is exceeded, then a specific action is taken as soon as the event occurs.</p>
Queuing	<p>Queuing is used to prevent traffic congestion. Traffic is sent to specific queues for servicing and scheduling based upon bandwidth allocation. Traffic is then scheduled or sent out through the port.</p>
Bandwidth	<p>Bandwidth allocation determines the available capacity for traffic that is subject to QoS policies.</p>
Trust	<p>Trust enables traffic to pass through the device, and the Differentiated Services Code Point (DSCP), precedence, or CoS values coming in from the end points are retained in the absence of any explicit policy configuration.</p>
Class of Service (CoS)	<p>CoS refers to the three bits in an 802.1Q header that are used to indicate the priority of the Ethernet frame as it passes through a switched network. The CoS bits in the 802.1Q header are commonly referred to as the 802.1p bits. 802.1Q is discarded prior to frame encapsulation in a VXLAN header, where CoS value is not present in VXLAN tunnel. To maintain QoS when a packet enters the VXLAN tunnel, the type of service (ToS) and CoS values map to each other.</p>
Differentiated Services Code Point (DSCP)	<p>The first six bits of the ToS byte in the IP header. DSCP is only present in an IP packet.</p>
Explicit Congestion Notification (ECN)	<p>The last two bits of the ToS byte in the IP header. ECN is only present in an IP packet.</p>

How to Configure BGP EVPN VXLAN QoS

QoS features in a BGP EVPN VXLAN fabric are enabled through the Modular QoS CLI (MQC). The MQC is a CLI structure that allows you to create traffic policies and attach these policies to interfaces. A traffic policy contains a traffic class and one or more QoS features. A traffic class is used to classify traffic, while the QoS features in the traffic policy determine how to treat the classified traffic.

MQC uses the following three steps to configure QoS:

1. Define a traffic class with the **class-map** command.

2. Create a traffic policy by associating the traffic class with one or more QoS features, using the **policy-map** command.
3. Attach the traffic policy to the interface with the **service-policy** command.

For more information about configuring QoS, see the *Quality of Service Configuration Guide, Cisco IOS XE* for the device.



Note In a BGP EVPN VXLAN fabric, QoS policies are applied on the physical interface of the VTEPs in the underlay. The NVE interface of the VTEPs support only the BUM rate limiter policies. For a sample configuration of BUM Traffic Rate Limiting, see "Configuring EVPN VXLAN Layer 2 Overlay Network" chapter in the *BGP EVPN VXLAN Configuration Guide* for the relevant release.

BGP EVPN VXLAN QoS Modes

Layer 2 or Layer 3 traffic enters a BGP EVPN VXLAN fabric at its edge node, which is also known as the ingress VTEP. The ingress VTEP encapsulates the original packet in a VXLAN header. The VXLAN header of the encapsulated packet is assigned a DSCP value based on QoS rules. The Spines and any other intermediate Layer 3 network devices are transit systems interconnecting the ingress and egress VTEPs. These intermediate devices use the outer header of the encapsulated packet to perform classification, marking, and policing. The encapsulated packet exits the fabric through the egress VTEP. At the egress VTEP, the VXLAN encapsulated packets are decapsulated and Ethernet packets are sent to the destination. The DSCP value of the inner header (decapsulated packet header) is used to perform classification, marking, and policing.

EVPN VXLAN QoS is supported in two modes:

Ingress VTEP: Uniform Mode

When a packet enters the fabric, if the ingress interface has a QoS policy that first marks the original packet and then the packet is encapsulated, both the inner and outer headers of the packet have the same updated DSCP value.

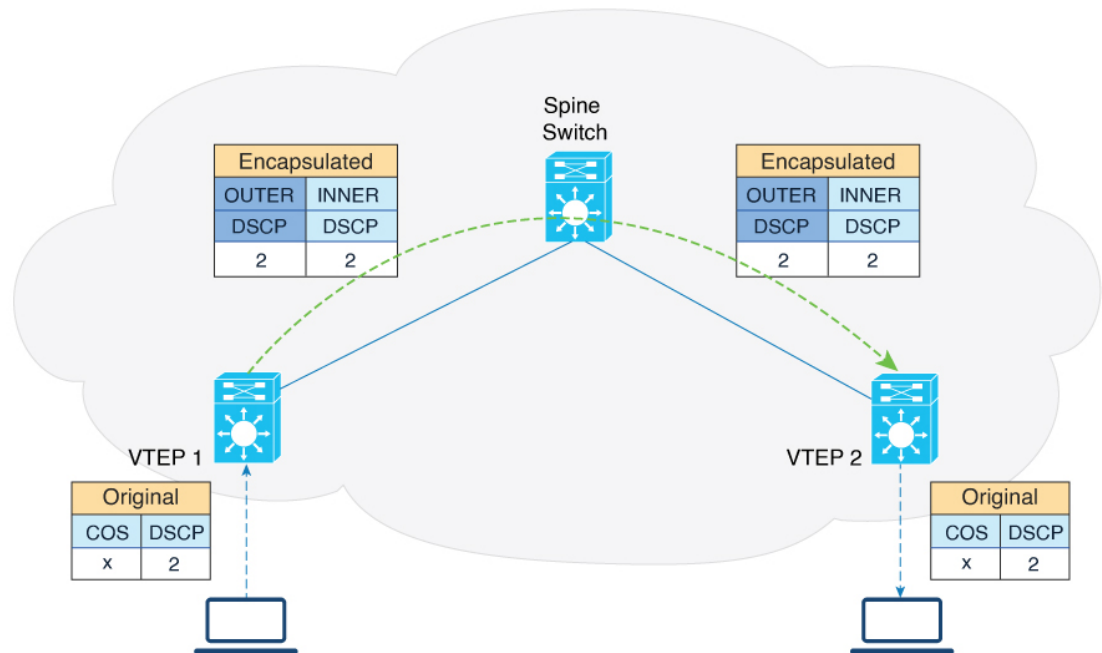
Uniform mode is the default mode of QoS operation.

Egress VTEP: Pipe Mode

As the encapsulated packet traverses through the fabric, if an intermediate node has a QoS policy that alters the outer DSCP value, the outer header is removed through decapsulation at egress. The DSCP value from the outer header is not propagated to the original packet. In this way, the original ToS value is preserved at the end of the VXLAN tunnel.

BGP EVPN VXLAN Traffic without QoS Policy

Figure 1: EVPN VXLAN Traffic without a QoS Policy



Encapsulation at Ingress VTEP

1. Traffic entering the edge of EVPN VXLAN fabric is received by the ingress VTEP (Leaf switch VTEP1) on its access interface.
2. Ingress VTEP encapsulates the traffic packets with a VXLAN header. During the VXLAN encapsulation, the DSCP value from inner header is copied to the outer header.

Transportation of a packet through the VXLAN Tunnel


As the VXLAN encapsulated packet traverses through the EVPN VXLAN fabric, the intermediate switches use the outer header to perform classification, marking, and policing.

Decapsulation at the Egress VTEP

1. When a VXLAN-encapsulated packet reaches the NVE interface of the egress VTEP (Leaf switch VTEP2), the switch removes the VXLAN header from the packet.
2. It uses the DSCP value of the inner header to perform classification, marking, and policing.
3. The DSCP value from the outer header is not propagated into the original packet header if the outer DSCP value was changed by any intermediate switch.
4. If a packet exits on a trunk port, the COS value for the packet is derived from the DSCP-to-COS table map.

BGP EVPN VXLAN QoS Marking with an Input Policy


Figure 2: EVPN QoS Marking with an Input Policy



1. The ingress VTEP VXLAN-encapsulates the packet and marks both the outer and inner header. During the VXLAN encapsulation, the IP packet's DSCP value is modified based on the input policy.
2. At egress, the outer VXLAN header is removed and the outer header DSCP values are not propagated to the decapsulated packets.

BGP EVPN VXLAN QoS Marking with Input and Output Policies

Figure 3: EVPN QoS Marking with an Input and an Output Policy



1. At ingress, the VTEP encapsulates the packet and marks both the outer and inner header. During the VXLAN encapsulation, the original packet's DSCP value is modified based on the input policy.
2. At egress, the outer VXLAN header is removed and the outer header DSCP or COS values are marked based on the egress QoS policy.