Configuring VXLAN QoS

This chapter contains these sections:

Information About VXLAN QoS

VXLAN QoS enables you to provide Quality of Service (QoS) capabilities to traffic that is tunneled in VXLAN.

Traffic in the VXLAN overlay can be assigned to different QoS properties:

- Classification traffic to assign different properties.
- Including traffic marking with different priorities.
- Queuing traffic to enable priority for the protected traffic.
- Policing for misbehaving traffic.
- Shaping for traffic that limits speed per interface.
- Properties traffic sensitive to traffic drops.

Note QoS allows you to classify the network traffic, police and prioritize the traffic flow, and provide congestion avoidance. For more information about QoS, see the Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide, Release 7.x.

Note QoS allows you to classify the network traffic, police and prioritize the traffic flow, and provide congestion avoidance. For more information about QoS, see the Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide, Release 9.2(x).

This section contains the following topics:

VXLAN QoS Terminology

This section defines VXLAN QoS terminology.

Table 1: VXLAN QoS Terminology

Term	Definition
Frames	Carries traffic at Layer 2. Layer 2 frames carry Layer 3 packets.
Packets	Carries traffic at Layer 3.
VXLAN packet	Carries original frame, encapsulated in VXLAN IP/UDP header.
Original frame	A Layer 2 or Layer 2 frame that carries the Layer 3 packet before encapsulation in a VXLAN header.
Decapsulated frame	A Layer 2 or a Layer 2 frame that carries a Layer 3 packet after the VXLAN header is decapsulated.
Ingress VTEP	The point where traffic is encapsulated in the VXLAN header and enters the VXLAN tunnel.
Egress VTEP	The point where traffic is decapsulated from the VXLAN header and exits the VXLAN tunnel.
Class of Service (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.
IP precedence	The 3 most significant bits of the ToS byte in the IP header.
Differentiated Services Code Point (DSCP)	The first six bits of the ToS byte in the IP header. DSCP is only present in an IP packet.
Explicit Congestion Notification (ECN)	The last two bits of the ToS byte in the IP header. ECN is only present in an IP packet.

Term	Definition
QoS tags	Prioritization values carried in Layer 3 packets and Layer 2 frames. A Layer 2 CoS label can have a value ranging between zero for low priority and seven for high priority. A Layer 3 IP precedence label can have a value ranging between zero for low priority and seven for high priority. IP precedence values are defined by the three most significant bits of the 1-byte ToS byte. A Layer 3 DSCP label can have a value between 0 and 63. DSCP values are defined by the six most significant bits of the 1-byte IP ToS field.
Classification	The process used for selecting traffic for QoS
Marking	The process of setting: a Layer 2 COS value in a frame, Layer 3 DSCP value in a packet, and Layer 3 ECN value in a packet. Marking is also the process of choosing different values for the CoS, DSCP, ECN field to mark packets so that they have the priority that they require during periods of congestion.
Policing	Limiting bandwidth used by a flow of traffic. Policing can mark or drop traffic.
MQC	The Cisco Modular QoS command line interface (MQC) framework, which is a modular and highly extensible framework for deploying QoS.

VXLAN QoS Features

The following topics describe the VXLAN QoS features that are supported in a VXLAN network:

Trust Boundaries

The trust boundary forms a perimeter on your network. Your network trusts (and does not override) the markings on your switch. The existing ToS values are trusted when received on in the VXLAN fabric.

Classification

You use classification to partition traffic into classes. You classify the traffic based on the port characteristics or the packet header fields that include IP precedence, differentiated services code point (DSCP), Layer 3 to Layer 4 parameters, and the packet length.

The values used to classify traffic are called match criteria. When you define a traffic class, you can specify multiple match criteria, you can choose to not match on a particular criterion, or you can determine the traffic class by matching any or all criteria.

Traffic that fails to match any class is assigned to a default class of traffic called class-default.

Marking

Marking is the setting of QoS information that is related to a packet. Packet marking allows you to partition your network into multiple priority levels or classes of service. You can set the value of a standard QoS field for COS, IP precedence, and DSCP. You can also set the QoS field for internal labels (such as QoS groups) that can be used in subsequent actions. Marking QoS groups is used to identify the traffic type for queuing and scheduling traffic.

Policing

Policing causes traffic that exceeds the configured rate to be discarded or marked down to a higher drop precedence.

Single-rate policers monitor the specified committed information rate (CIR) of traffic. Dual-rate policers monitor both CIR and peak information rate (PIR) of traffic.

Queuing and Scheduling

The queuing and scheduling process allows you to control the queue usage and the bandwidth that is allocated to traffic classes. You can then achieve the desired trade-off between throughput and latency.

You can limit the size of the queues for a particular class of traffic by applying either static or dynamic limits.

You can apply weighted random early detection (WRED) to a class of traffic, which allows packets to be dropped based on the QoS group. The WRED algorithm allows you to perform proactive queue management to avoid traffic congestion.

ECN can be enabled along with WRED on a particular class of traffic to mark the congestion state instead of dropping the packets. ECN marking in the VXLAN tunnel is performed in the outer header, and at the Egress VTEP is copied to decapsulated frame.

Traffic Shaping

You can shape traffic by imposing a maximum data rate on a class of traffic so that excess packets are retained in a queue to smooth (constrain) the output rate. In addition, minimum bandwidth shaping can be configured to provide a minimum guaranteed bandwidth for a class of traffic.

Traffic shaping regulates and smooths out the packet flow by imposing a maximum traffic rate for each port's egress queue. Packets that exceed the threshold are placed in the queue and are transmitted later. Traffic shaping is similar to Traffic Policing, but the packets are not dropped. Because packets are buffered, traffic shaping minimizes packet loss (based on the queue length), which provides better traffic behavior for TCP traffic.

By using traffic shaping, you can control the following:

- Access to available bandwidth.
- Ensure that traffic conforms to the policies established for it.
- Regulate the flow of traffic to avoid congestion that can occur when the egress traffic exceeds the access speed of its remote, target interface.

For example, you can control access to the bandwidth when the policy dictates that the rate of a given interface must not, on average, exceed a certain rate. Despite the access rate exceeding the speed.

Network QoS

The network QoS policy defines the characteristics of each CoS value, which are applicable network wide across switches. With a network QoS policy, you can configure the following:

- Pause behavior—You can decide whether a CoS requires the lossless behavior which is provided by using a priority flow control (PFC) mechanism that prevents packet loss during congestion) or not. You can configure drop (frames with this CoS value can be dropped) and no drop (frames with this CoS value cannot be dropped). For the drop and no drop configuration, you must also enable PFC per port. For more information about PFC, see “Configuring Priority Flow Control".

Pause behavior can be achieved in the VXLAN tunnel for a specific queue-group.

VXLAN Priority Tunneling

In the VXLAN tunnel, DSCP values in the outer header are used to provide QoS transparency in end-to-end of the tunnel. The outer header DSCP value is derived from the DSCP value with Layer 3 packets or the CoS value for Layer 2 frames. At the VXLAN tunnel egress point, the priority of the decapsulated traffic is chosen based on the mode. For more information, see Decapsulated Packet Priority Selection, on page 9.

MQC CLI

All available QoS features for VXLAN QoS are managed from the modular QoS command-line interface (CLI). The Modular QoS CLI (MQC) allows you to define traffic classes (class maps), create and configure traffic policies (policy maps), and perform actions that are defined in the policy maps to interface (service policy).

VXLAN QoS Topology and Roles

This section describes the roles of network devices in implementing VXLAN QoS.

Figure 1: VXLAN Network

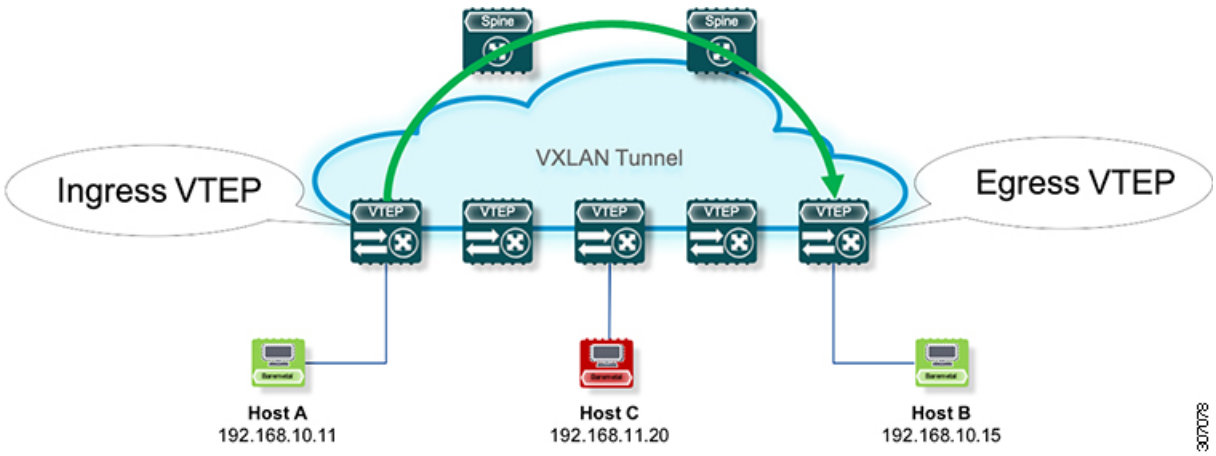

The network is bidirectional, but in the previous image, traffic is moving left to right.

In the VXLAN network, points of interest are ingress VTEPs where the original traffic is encapsulated in a VXLAN header. Spines are transporting hops that connect ingress and egress VTEPs. An egress VTEP is the point where VXLAN encapsulated traffic is decapsulated and egresses the VTEP as classical Ethernet traffic.

Note Ingress and egress VTEPs are the boundary between the VXLAN tunnel and the IP network.

This section contains the following topics:

Ingress VTEP and Encapsulation in the VXLAN Tunnel

At the ingress VTEP, the VTEP processes packets as follows:

Procedure

Step 1 Layer 2 or Layer 3 traffic enters the edge of the VXLAN network.

Step 2 The switch receives the traffic from the input interface and uses the 802.1p bits or the DSCP value to perform any classification, marking, and policing. It also derives the outer DSCP value in the VXLAN header. For classification of incoming IP packets, the input service policy can also use access control lists (ACLs).

Step 3 For each incoming packet, the switch performs a lookup of the IP address to determine the next hop.

Step 4 The packet is encapsulated in the VXLAN header. The encapsulated packet's VXLAN header is assigned a DSCP value that is based on QoS rules.

Step 5 The switch forwards the encapsulated packets to the appropriate output interface for processing.

Step 6 The encapsulated packets, marked by the DSCP value, are sent to the VXLAN tunnel output interface.

Transport Through the VXLAN Tunnel

In the transport through a VXLAN tunnel, the switch processes the VXLAN packets as follows:

Procedure

Step 1 The VXLAN encapsulated packets are received on an input interface of a transport switch. The switch uses the outer header to perform classification, marking, and policing.

Step 2 The switch performs a lookup on the IP address in the outer header to determine the next hop.

Step 3 The switch forwards the encapsulated packets to the appropriate output interface for processing.

Step 4 VXLAN sends encapsulated packets through the output interface.

Egress VTEP and Decapsulation of the VXLAN Tunnel

At the egress VTEP boundary of the VXLAN tunnel, the VTEP processes packets as follows:

Procedure

Step 1 Packets encapsulated in VXLAN are received at the NVE interface of an egress VTEP, where the switch uses the inner header DSCP value to perform classification, marking, and policing.

Step 2 The switch removes the VXLAN header from the packet, and does a lookup that is based on the decapsulated packet's headers.

Step 3 The switch forwards the decapsulated packets to the appropriate output interface for processing.

Step 4 Before the packet is sent out, a DSCP value is assigned to a Layer 3 packet based on the decapsulation priority or based on marking Layer 2 frames.

Step 5 The decapsulated packets are sent through the outgoing interface to the IP network.

Classification at the Ingress VTEP, Spine, and Egress VTEP

This section includes the following topics:

IP to VXLAN

At the ingress VTEP, the ingress point of the VXLAN tunnel, traffic is encapsulated in the VXLAN header. Traffic on an ingress VTEP is classified based on the priority in the original header. Classification can be performed by matching the CoS, DSCP, and IP precedence values or by matching traffic with the ACL based on the original frame data.

When traffic is encapsulated in the VXLAN, the Layer 3 packet's DSCP value is copied from the original header to the outer header of the VXLAN encapsulated packet. This behavior is illustrated in the following figure:

Figure 2: Copy of Priority from Layer-3 Packet to VXLAN Outer Header

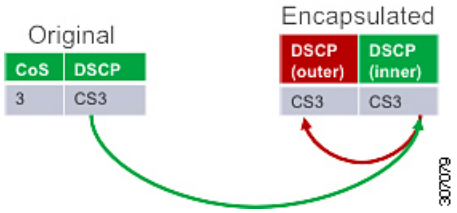

For Layer 2 frames without the IP header, the DSCP value of the outer header is derived from the CoS-to-DSCP mapping present in the hardware illustrated in Default Settings for VXLAN QoS, on page 14. In this way, the original QoS attributes are preserved in the VXLAN tunnel. This behavior is illustrated in the following figure:

Figure 3: Copy of Priority from Layer-2 Frame to VXLAN Outer Header

A Layer 2 frame, does not have a DSCP value present because the IP header is not present in the frame. After a Layer 2 frame is encapsulated, the original CoS value is not preserved in the VXLAN tunnel.

IP to VXLAN with Outer DSCP

Beginning with Cisco NX-OS Release 10.4(1)F, the policy with set outer DSCP action can be applied to the access interface in the ingress direction.

When traffic is encapsulated in the VXLAN, for Layer-3 packets, DSCP value from the original packet is copied to the inner header and the user configured DSCP value is set in the outer header of the VXLAN encapsulated packet. This behavior is illustrated in the following figure:

Figure 4: VXLAN Outer DSCP Value Applied from Set Configuration

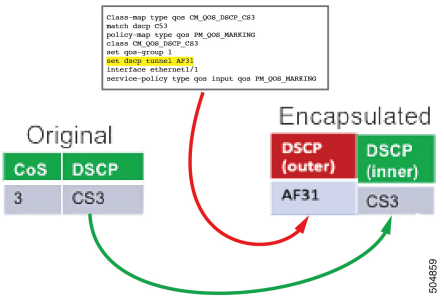

Inside the VXLAN Tunnel

Inside the VXLAN tunnel, traffic classification is based on the outer header DSCP value. Classification can be done matching the DCSP value or using ACLs for classification.

If VXLAN encapsulated traffic is crossing the trust boundary, marking can be changed in the packet to match QoS behavior in the tunnel. Marking can be performed inside of the VXLAN tunnel, where a new DSCP value is applied only on the outer header. The new DSCP value can influence different QoS behaviors inside the VXLAN tunnel. The original DSCP value is preserved in the inner header.

Figure 5: Marking Inside of the VXLAN Tunnel

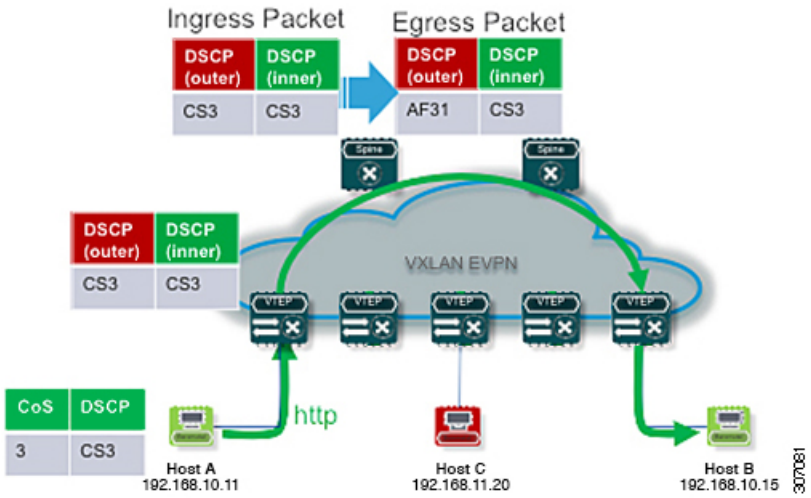

VXLAN to IP

Classification at the egress VTEP is performed for traffic leaving the VXLAN tunnel. For classification at the egress VTEP, the inner header and outer DSCP values are used. The inner or outer DSCP value is used for priority-based classification. Classification can be performed using ACLs.

Classification is performed on the NVE interface for all VXLAN tunneled traffic.

Marking and policing can be performed on the NVE interface for tunneled traffic. If marking is configured, newly marked values are present in the decapsulated packet. Because the original CoS value is not preserved in the encapsulated packet, marking can be performed for decapsulated packets for any devices that expect an 802.1p field for QoS in the rest of the network.

Decapsulated Packet Priority Selection

At the egress VTEP, the VXLAN header is removed from the packet and the decapsulated packet egresses the switch with the DSCP value. The switch assigns the DSCP value of the decapsulated packet based on two modes:

- Uniform mode – the DSCP value from the outer header of the VXLAN packet is copied to the decapsulated packet. Any change of the DSCP value in the VXLAN tunnel is preserved and present in the decapsulated packet. Uniform mode is the default mode of decapsulated packet priority selection.

Figure 6: Uniform Mode Outer DSCP Value is Copied to Decapsulated Packet DSCP Value for a Layer-3 Packet

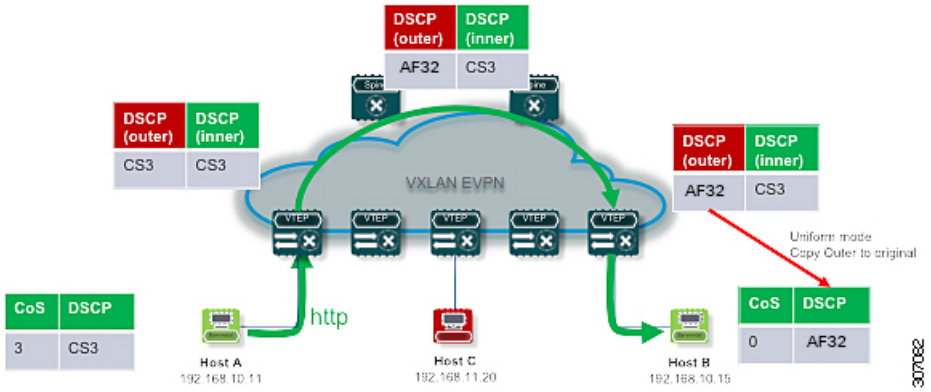

- Pipe mode – the original DSCP value is preserved at the VXLAN tunnel end. At the egress VTEP, the system copies the inner DSCP value to the decapsulated packet DSCP value. In this way, the original DSCP value is preserved at the end of the VXLAN tunnel.

Figure 7: Pipe Mode Inner DSCP Value is Copied to Decapsulated Packet DSCP Value for Layer-3 Packet

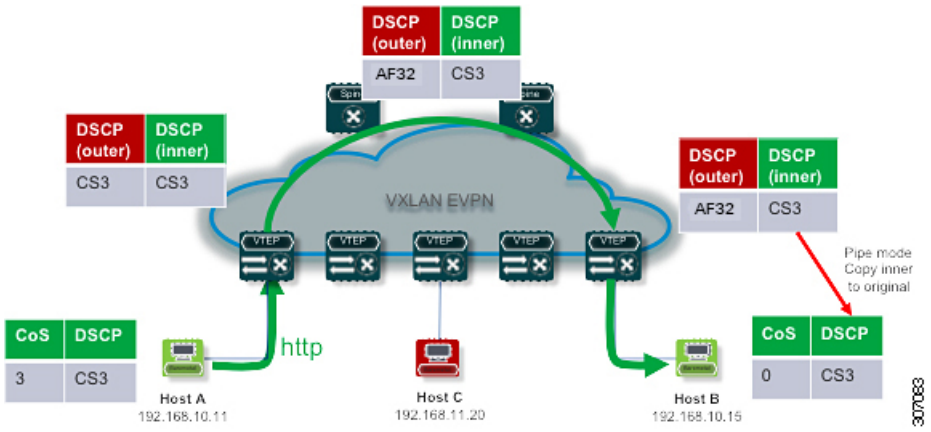

CoS Preservation

Beginning with Cisco NX-OS Release 10.4(1)F, the **default-vxlan-in-tnl-dscp-policy** QoS policy-map template is added to provide CoS preservation for non-IP packets.

When this template is enabled on the NVE interface, the switch performs a match on the outer DSCP of the VXLAN packet and rewrites the CoS in the decapsulated ethernet packet on the egress VTEP based on a fixed outer DSCP to CoS mapping.

The following table lists the default Outer DSCP-to-CoS mapping in the Egress VTEP for Layer 2 frames:

Table 2: Default Outer DSCP-to-CoS Mapping

DSCP of Outer VXLAN Header	CoS of Original Layer 2 Frame
0	0
8	1
16	2
26	3
32	4
46	5
48	6
56	7

Figure 8: Non-IP CoS Value Restored in Decapsulated Packet

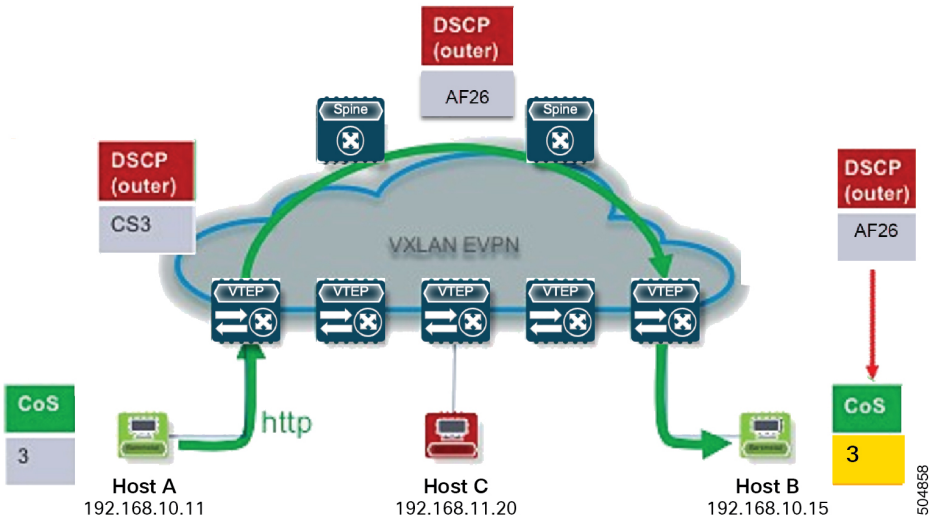

Guidelines and Limitations for VXLAN QoS

Note The QoS policy must be configured end-to-end for this feature to work as designed.

VXLAN QoS has the following configuration guidelines and limitations:

- Cisco Nexus 9364C, 9300-EX, and 9300-FX/FX2/FX3 platform switches and Cisco Nexus 9500 platform switches with EX/FX or -R/RX line cards support VXLAN QoS.
- Beginning with Cisco NX-OS Release 9.3(3), Cisco Nexus 9300-GX platform switches support VXLAN QoS in default mode.
- Beginning with Cisco NX-OS Release 10.2(3)F, VXLAN QoS in default mode is supported on Cisco Nexus 9300-GX2 platform switches.
- Beginning with Cisco NX-OS Release 10.4(1)F, VXLAN QoS in default mode is supported on Cisco Nexus 9332D-H2R switches.
- Beginning with Cisco NX-OS Release 10.4(2)F, VXLAN QoS in default mode is supported on Cisco Nexus 93400LD-H1 switches.
- Beginning with Cisco NX-OS Release 10.4(3)F, VXLAN QoS in default mode is supported on Cisco Nexus 9364C-H1 switches.
- The following features are supported on Cisco Nexus 9504 and 9508 platform switches with -R/RX line cards:
  - Physical interface level queuing should work as normal L2/L3 queuing/QoS
  - IPv4 bridged case works in terms of copying inner ToS to outer VXLAN ToS

- The following features are not supported on Cisco Nexus 9504 and 9508 platform switches with -R and -RX line cards:
  - Policies on the NVE interface
  - IPv6 type of service (ToS) from inner to VXLAN outer copying
  - IPv4 routed cases for QoS. ToS from inner is not copied to outer VXLAN header
- For Cisco Nexus 9504 and 9508 platform switches with -RX line cards, the default mode is pipe for VXLAN decapsulation (inner packet DSCP not modified based on outer IP header DSCP value). This is a difference in behavior from other line cards types. If -RX line cards and other line cards are used in the same network, the **qos-mode pipe** command can be used in switches where non-RX line cards are present in order to have the same behavior. For details on the configuration command, see Configuring Type QoS on the Egress VTEP, on page 15.
- VXLAN QoS is supported in the EVPN fabric.
- The original IEEE 802.1Q header is not preserved in the VXLAN tunnel. The CoS value is not present in the inner header of the VXLAN-encapsulated packet.
- Statistics (counters) are present for the NVE interface.
- Egress policing is not supported on outgoing interface (uplink connecting to spine) of the encap (ingress) VXLAN VTEP.
- In a vPC, configure the change of the decapsulated packet priority selection on both peers.
- The service policy on an NVE interface can attach only in the input direction.
- If DSCP marking is present on the NVE interface, traffic to the BUD node preserves marking in the inner and outer headers. If a marking action is configured on the NVE interface, BUM traffic is marked with a new DSCP value on Cisco Nexus 9364C and 9300-EX platform switches.
- A classification policy applied to an NVE interface applies only on VXLAN-encapsulated traffic. For all other traffic, the classification policy must be applied on the incoming interface.
- To mark the decapsulated packet with a CoS value, a marking policy must be attached to the NVE interface to mark the CoS value to packets where the VLAN header is present.
- The following guidelines and limitations apply to VXLAN QoS configuration on the DCI handoff node:
  - Beginning with Cisco NX-OS Release 9.3(5), Cisco Nexus 9300-GX platform switches support VXLAN QoS configuration on the DCI handoff node.
  - Beginning with Cisco NX-OS Release 10.2(3)F, Cisco Nexus 9300-GX2 platform switches support VXLAN QoS configuration on the DCI handoff node.
  - Beginning with Cisco NX-OS Release 10.4(1)F, Cisco Nexus 9332D-H2R switches support VXLAN QoS configuration on the DCI handoff node.
  - Beginning with Cisco NX-OS Release 10.4(2)F, Cisco Nexus 93400LD-H1 switches support VXLAN QoS configuration on the DCI handoff node.
  - Beginning with Cisco NX-OS Release 10.4(3)F, Cisco Nexus 9364C-H1 switches support VXLAN QoS configuration on the DCI handoff node.
  - VXLAN QoS configuration on the DCI handoff node does not support end-to-end priority flow control (PFC) for Cisco Nexus 9336C-FX2, 93240YC-FX2, and 9300-GX platform switches.

- Microburst, dynamic packet prioritization (DPP), and approximate fair-drop (AFD) are supported on VXLAN-encapsulated packets.

- The following guidelines and limitations apply to the outer DSCP based VXLAN QoS Policy feature:
  - Beginning with Cisco NX-OS Release 10.4(1)F, the outer DSCP based VXLAN QoS Policy feature is supported on Cisco Nexus 9300-FX2/FX3/GX/GX2 platform switches and 9500 switch with N9K-X9716D-GX line card.
  - Beginning with Cisco NX-OS Release 10.4(2)F, the outer DSCP based VXLAN QoS Policy feature is supported on Cisco Nexus 9332D-H2R, and 93400LD-H1 switches.
  - Beginning with Cisco NX-OS Release 10.4(3)F, the outer DSCP based VXLAN QoS Policy feature is supported on Cisco Nexus 9364C-H1 switches.
  - In VXLAN QoS policy, the **match dscp tunnel** command can only be applied on the NVE interface and in ingress direction.
  - In VXLAN QoS Policy, both inner and outer DSCP match rules is not supported. However, the match criteria like **ip access-lists** or **mac access-list** in the same policy applied on the NVE interface would always match on the inner header.
  - For non-IP packets, the outer header QoS policy under the NVE interface only supports L2 rewrite and traffic class assignment or outgoing queue. Action like policer is not supported.
  - In VXLAN QoS policy, the **match dscp tunnel** command on the NVE interface performs match on the VXLAN packets destined to the current VTEP, where tunnel termination happens, and packets are decapsulated.
  - In VXLAN QoS policy, the **match dscp tunnel** command does not support non-IP packets, due to this the CoS preservation will not work on IPv6 underlay.
  - In VXLAN QoS policy, the **set dscp tunnel** command does not support non-IP packets. In case of non-IP packets, the outer DSCP value is applied based on the default CoS to DSCP mapping information on the switch.
  - In VXLAN QoS policy, the **set dscp tunnel** command cannot be applied to the NVE interface, as this command is applicable to encapsulation packets.
  - If the **set dscp tunnel** command is applied on the ingress VTEP of VXLAN multisite, the outer DSCP value could be replaced with the inner DSCP if pipe mode is configured on the border gateway. We recommend to configure uniform mode on the border gateway to carry the new outer DSCP header to the remote site.
  - The outer DSCP based VXLAN QoS Policy feature is not supported on the VXLAN multisite deployments.

- The following limitations apply to the VXLAN QoS policies when using a Border Gateway (BGW) Spine:
  - If QoS policies are needed for intra-site BUM traffic for VNI with multicast underlay, and that multicast underlay group is also owned by a VNI defined on the BGW Spine, then the QoS policy must be applied to the NVE interface. QoS policies applied to fabric interfaces will not modify these flows since the NVE interface acts as an incoming interface.
  - If QoS policies are needed for intra-site BUM traffic for VNI with multicast underlay, and that multicast group is not owned by a VNI defined on the BGW Spine, then the QoS policy must be

applied to a fabric interface. QoS policies applied to the NVE interface will not modify these flows since the NVE is not considered an incoming interface.

- If the NVE interface of the BGW Spine owns a multicast group used for BUM traffic within the local fabric, QoS policies cannot be applied to both the fabric interfaces and NVE interface to differentiate treatment of intra-site and inter-site flows for that multicast group.

- Beginning with Cisco NX-OS Release 10.4(3)F, Cisco Nexus 9808/9804 switches with X9836DM-A and X98900CD-A line cards support VXLAN QoS policies when using a BGW spine with the following limitations:
  - Physical ingress QoS policy and system-level QoS policy are supported.
  - QoS policy on NVE is not supported.
  - Explicit Congestion Notification (ECN) or ECN-Capable Transport (ECT) marking is not retained.

Default Settings for VXLAN QoS

The following table lists the default CoS-to-DSCP mapping in the ingress VTEP for Layer 2 frames:

Table 3: Default CoS-to-DSCP Mapping

CoS of Original Layer 2 Frame	DSCP of Outer VXLAN Header
0	0
1	8
2	16
3	26
4	32
5	46
6	48
7	56

Configuring VXLAN QoS

Configuration of VXLAN QoS is done using the MQC model. The same configuration that is used for the QoS configuration applies to VXLAN QoS. For more information about configuring QoS, see the Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide, Release 9.2(x).

VXLAN QoS introduces a new service-policy attachment point which is NVE – Network Virtual Interface. At the egress VTEP, the NVE interface is the point where traffic is decapsulated. To account for all VXLAN traffic, the service policy must be attached to an NVE interface.

The next section describes the configuration of the classification at the egress VTEP, and **service-policy type qos** attachment to an NVE interface.

Configuring Type QoS on the Egress VTEP

Configuration of VXLAN QoS is done by using the MQC model. The same configuration is used for QoS configuration for VXLAN QoS. For more information about configuring QoS, see the Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide, Release 9.2(x).

VXLAN QoS introduces a new service-policy attachment point which is the Network Virtual Interface (NVE). At the egress VTEP, the NVE interface points where traffic is decapsulated. To account for all VXLAN traffic, the service policy must be attached to an NVE interface.

This procedure describes the configuration of classification at the egress VTEP, and **service-policy type qos** attachment to an NVE interface.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enters global configuration mode.
Step 2	**[no] class-map [type [qos]] | [match-all] | [match-any]** *class-map-name* **Example:** `switch(config)# class-map type qos class1`	Creates or accesses the class map *class--map-name* and enters **class-map** mode. The *class--map-name* argument can contain alphabetic, hyphen, or underscore characters, and can be up to 40 characters. (match-any is the default when the **no** option is selected and multiple match statements are entered.)
Step 3	**[no] match [access-group | cos | dscp [tunnel] | precedence]** *{name | 0-7 | 0-63 | 0-7}* **Example:** `switch(config-cmap-qos)# match dscp tunnel 26`	Configures the traffic class by matching packets based on access-list, **cos** value, **dscp** values, or IP **precedence** value Beginning with Cisco NX-OS Release 10.4(1)F, the **tunnel** option is provided to match the DSCP value on the outer VXLAN header of the ingress packet. **Note** The **match dscp tunnel** command will be used in an ingress service policy that will be applied to the NVE interface on the egress VTEP.
Step 4	**[no] policy-map type qos** *policy-map-name* **Example:** `switch(config-cmap-qos)# policy-map type qos policy`	Creates or accesses the policy map that is named *policy-map-name* and then enters policy-map mode. The policy-map name can contain alphabetic, hyphen, or underscore characters, is case sensitive, and can be up to 40 characters.
Step 5	**[no] class** *class-name* **Example:** `switch(config-pmap-qos)# class class1`	Creates a reference to class-name and enters policy-map class configuration mode. The class is added to the end of the policy map unless insert-before is used to specify the

	Command or Action	Purpose
		class to insert before. Use the class-default keyword to select all traffic that is not currently matched by classes in the policy map.
Step 6	**[no] set qos-group** *qos-group-value* **Example:** `switch(config-pmap-c-qos)# set qos-group 1`	Sets the QoS group value to *qos-group-value*. The value can range from 1 through 126. The **qos-group** is referenced in type queuing and type network-qos as matching criteria.
Step 7	**exit** **Example:** `switch(config-pmap-c-qos)# exit`	Exits class-map mode.
Step 8	**[no] interface nve** *nve-interface-number* **Example:** `switch(config)# interface nve 1`	Enters interface mode to configure the NVE interface.
Step 9	**[no] service-policy type qos input** *policy-map-name* **Example:** `switch(config-if-nve)# service-policy type qos input policy`	Adds a service-policy *policy-map-name* to the interface in the input direction. You can attach only one input policy to an NVE interface.
Step 10	(Optional) **[no] qos-mode [pipe]** **Example:** `switch(config-if-nve)# qos-mode pipe`	Selecting decapsulated packet priority selection and using pipe mode. Entering the **no** form of this command negates pipe mode and defaults to uniform mode.

Setting Outer DSCP on the Ingress VTEP

VXLAN QoS policy introduces a new outer DSCP set action for all VXLAN traffic, the service policy must be attached to the access (ingress) interface of the ingress VTEP.

SUMMARY STEPS

1. **configure terminal**
2. **[no] class-map [type qos] [match-all]| [match-any]** *class-map-name*
3. **[no] policy-map type qos** *policy-map-name*
4. **[no] class** *class-name*
5. **[no] set dscp [tunnel]** *dscp-val*

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enters global configuration mode.
Step 2	**[no] class-map [type qos] [match-all]| [match-any]** *class-map-name* **Example:** `switch(config)# class-map type qos class1`	Creates or accesses the class map *class--map-name* and enters **class-map** mode. The *class--map-name* argument can contain alphabetic, hyphen, or underscore characters, and can be up to 40 characters. (match-any is the default when the **no** option is selected and multiple match statements are entered.)
Step 3	**[no] policy-map type qos** *policy-map-name* **Example:** `switch(config-cmap-qos)# policy-map type qos policy`	Creates or accesses the policy map that is named *policy-map-name* and then enters policy-map mode. The policy-map name can contain alphabetic, hyphen, or underscore characters, is case sensitive, and can be up to 40 characters.
Step 4	**[no] class** *class-name* **Example:** `switch(config-pmap-qos)# class class1`	Creates a reference to class-name and enters policy-map class configuration mode. The class is added to the end of the policy map unless insert-before is used to specify the class to insert before. Use the class-default keyword to select all traffic that is not currently matched by classes in the policy map.
Step 5	**[no] set dscp [tunnel]** *dscp-val* **Example:** `switch(config-pmap-c-qos)# set dscp tunnel 32`	Sets the DSCP value in the outer VXLAN header of the Ingress packet.

Verifying the VXLAN QoS Configuration

Table 4: VXLAN QoS Verification Commands

Command	Purpose
show class map	Displays information about all configured class maps.
show policy-map	Displays information about all configured policy maps.
show running ipqos	Displays configured QoS configuration on the switch.

VXLAN QoS Configuration Examples

Ingress VTEP Classification and Marking

This example shows how to configure the **class-map type qos** command for classification matching traffic with an ACL. Enter the **policy-map type qos** command to put traffic in qos-group 1 and set the DSCP value. Enter the **service-policy type qos** command to attach to the ingress interface in the input direction to classify traffic matching the ACL.

```
access-list ACL_QOS_DSCP_CS3 permit ip any any eq 80

class-map type qos CM_QOS_DSCP_CS3
 match access-group name ACL_QOS_DSCP_CS3

policy-map type qos PM_QOS_MARKING
 class CM_QOS_DSCP_CS3
 set qos-group 1
 set dscp 24

interface ethernet1/1
 service-policy type qos input PM_QOS_MARKING
```

Transit Switch – Spine Classification

This example shows how to configure the **class-map type qos** command for classification matching DSCP 24 set on the ingress VTEP. Enter the **policy-map type qos** command to put traffic in qos-group 1. Enter the **service-policy type qos** command to attach to the ingress interface in the input direction to classify traffic matching criteria.

```
class-map type qos CM_QOS_DSCP_CS3
 match dscp 24

policy-map type qos PM_QOS_CLASS
 class CM_QOS_DSCP_CS3
 set qos-group 1

interface Ethernet 1/1
 service-policy type qos input PM_QOS_CLASS
```

Egress VTEP Classification and Marking

This example shows how to configure the **class-map type qos** command for classification matching traffic by DSCP value. Enter the **policy-map type qos** to place traffic in qos-group 1 and mark CoS value in outgoing frames. The **service-policy type qos** command is applied to the NVE interface in the input direction to classify traffic coming out of the VXLAN tunnel.

```
class-map type qos CM_QOS_DSCP_CS3
 match dscp 24

policy-map type qos PM_QOS_MARKING
 class CM_QOS_DSCP_CS3
 set qos-group 1
 set cos 3

interface nve 1
```

```
 service-policy type qos input PM_QOS_MARKING
```

Queuing

This example shows how to configure the **policy-map type queueing** command for traffic in qos-group 1. Assigning 50% of the available bandwidth to q1 mapped to qos-group 1 and attaching policy in the output direction to all ports using the **system qos** command.

```
policy-map type queuing PM_QUEUING
class type queuing c-out-8q-q7
      priority level 1
    class type queuing c-out-8q-q6
      bandwidth remaining percent 0
    class type queuing c-out-8q-q5
      bandwidth remaining percent 0
    class type queuing c-out-8q-q4
      bandwidth remaining percent 0
    class type queuing c-out-8q-q3
      bandwidth remaining percent 0
    class type queuing c-out-8q-q2
      bandwidth remaining percent 0
    class type queuing c-out-8q-q1
      bandwidth remaining percent 50
    class type queuing c-out-8q-q-default
      bandwidth remaining percent 50

system qos
 service-policy type queueing output PM_QUEUING
```

Preserve CoS Configuration

This example shows how to configure the CoS preservation on NVE interface:

```
interface nve 1
   service-policy type qos input default-vxlan-in-tnl-dscp-policy
```