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.

- 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.

Additional Information About VXLAN QoS

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 following guides:

- Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide, Release 7.x
- Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide, Release 9.2(x)

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

Classification partitions network traffic into classes based on port characteristics or packet header fields, including IP precedence, DSCP, Layer 3 to Layer 4 parameters, and packet length.

- Match criteria: The values used to classify traffic are called match criteria.
- Multiple match criteria: When defining a traffic class, you can specify multiple match criteria, choose not to match on a particular criterion, or determine the traffic class by matching any or all criteria.
- Default class: Traffic that fails to match any class is assigned to a default traffic class called `class-default`.

Marking

Marking sets QoS information associated with 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 Class of Service (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 identifies the traffic type for queuing and scheduling traffic.

- Partition the network into multiple priority levels or classes of service.
- Set the value of a standard QoS field for CoS, IP precedence, and DSCP.
- Set the QoS field for internal labels (such as QoS groups) for use in subsequent actions.

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.

How does traffic shaping work?

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.

Network QoS

The network QoS policy defines the characteristics of each CoS value, which are applicable network wide across switches.

- 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 10.

- DSCP values in the outer header provide QoS transparency.
- Outer header DSCP is derived from Layer 3 DSCP or Layer 2 CoS values.
- Priority of decapsulated traffic at egress is mode-dependent.

MQC 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).

- Define traffic classes (class maps)
- Create and configure traffic policies (policy maps)
- Perform actions that are defined in the policy maps to interface (service policy)

VXLAN QoS Topology and Roles

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.

- Ingress VTEP: Encapsulates original traffic in a VXLAN header.
- Spine: Transports traffic between ingress and egress VTEPs.

- Egress VTEP: Decapsulates VXLAN traffic and outputs as Ethernet.

VXLAN QoS Topology Reference

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

Key roles in the VXLAN QoS topology include:

- Ingress VTEP
- Spine
- Egress VTEP

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

The following figure shows the VXLAN network topology:

Figure 1: VXLAN Network

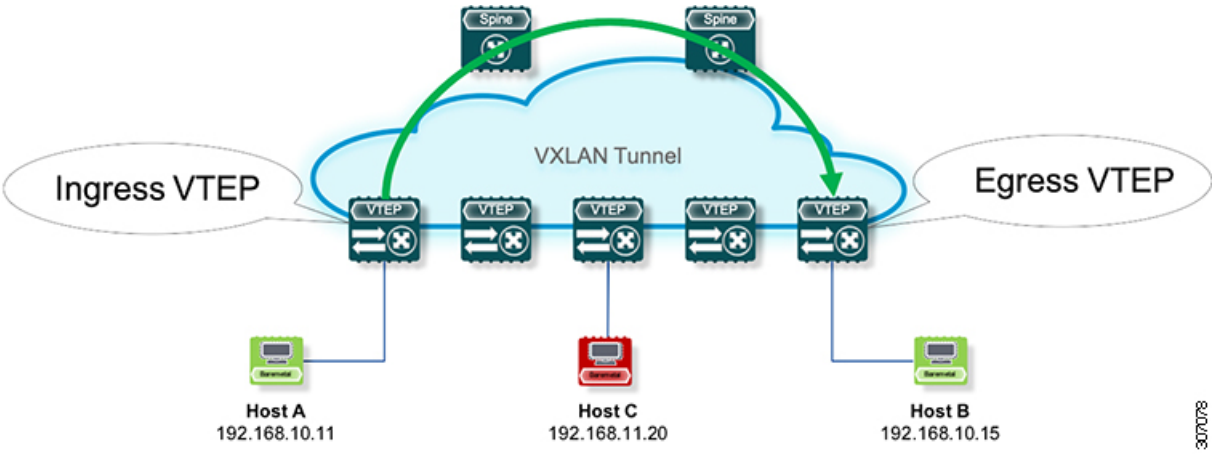

Ingress VTEP and Encapsulation in the VXLAN Tunnel

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

Summary

The ingress VTEP receives Layer 2 or Layer 3 traffic, classifies and marks packets based on QoS policies, determines the next hop, and encapsulates packets with VXLAN headers for forwarding through the VXLAN tunnel.

- Ingress VTEP: Receives and classifies incoming traffic, applies QoS markings, and encapsulates packets into VXLAN headers.
- Switch: Forwards encapsulated packets through the VXLAN tunnel based on routing and QoS parameters.

This process ensures that traffic entering the VXLAN fabric is properly encapsulated, marked, and forwarded according to QoS and routing policies.

Workflow

These stages describe how the ingress VTEP processes and encapsulates packets in the VXLAN tunnel.

1. Layer 2 or Layer 3 traffic enters the edge of the VXLAN network.
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).
3. For each incoming packet, the switch performs a lookup of the IP address to determine the next hop.
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.
5. The switch forwards the encapsulated packets to the appropriate output interface for processing.
6. The encapsulated packets, marked by the DSCP value, are sent to the VXLAN tunnel output interface.

Transporting VXLAN Packets

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

Summary

The switch receives VXLAN-encapsulated packets, performs header-based classification and forwarding decisions, and sends the packets through the appropriate output interface for transport across the VXLAN tunnel.

Workflow

These stages describe how VXLAN encapsulated packets are processed and forwarded through the tunnel.

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.
2. The switch performs a lookup on the IP address in the outer header to determine the next hop.
3. The switch forwards the encapsulated packets to the appropriate output interface for processing.
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:

Summary

The egress VTEP receives VXLAN-encapsulated packets, removes the VXLAN header, and performs forwarding decisions based on the inner headers. The switch applies QoS marking and policing, assigns DSCP values as required, and sends the decapsulated packets to the destination network through the appropriate output interface.

Workflow

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.
2. The switch removes the VXLAN header from the packet, and does a lookup that is based on the decapsulated packet's headers.
3. The switch forwards the decapsulated packets to the appropriate output interface for processing.
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.
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.

- Classification can be performed by matching the CoS, DSCP, and IP precedence values.
- Classification can also be performed by matching traffic with the ACL based on the original frame data.
- 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.

Preservation of QoS Attributes in VXLAN Encapsulation

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.

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 13. In this way, the original QoS attributes are preserved in the VXLAN tunnel.

This behavior is illustrated in the following figures:

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

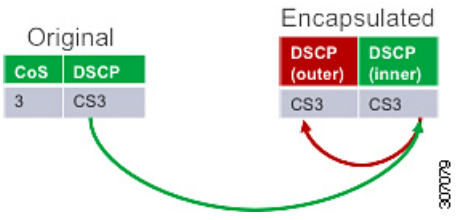

Figure 3: Copy of Priority from Layer-2 Frame to VXLAN Outer Header

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.

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.

Marking and QoS Behavior in VXLAN Tunnels

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.

Key points about marking and QoS behavior in VXLAN tunnels:

- Marking can be changed when crossing the trust boundary.
- A new DSCP value is applied only on the outer header.
- The original DSCP value is preserved in the inner header.

The following figure illustrates marking inside of the VXLAN tunnel:

Figure 4: Marking Inside of the VXLAN Tunnel

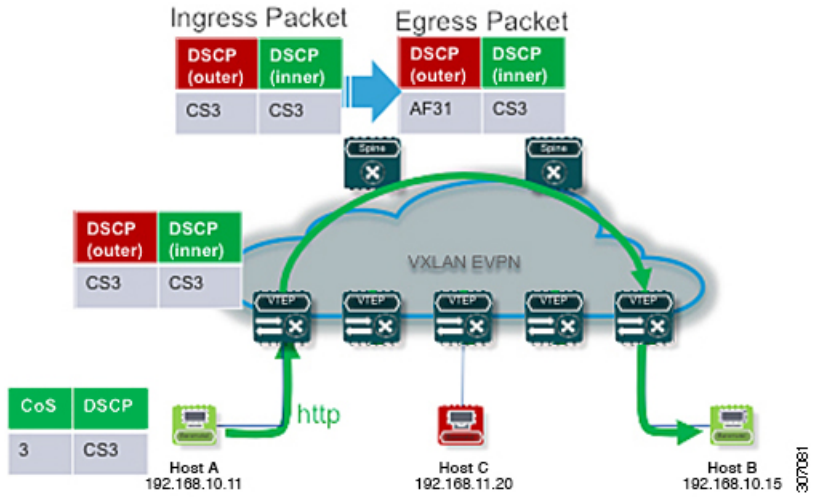

VXLAN to IP

Classification at the egress VTEP is performed for traffic leaving the VXLAN tunnel. For classification at the egress VTEP, the inner header values are used. The inner 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 5: Uniform Mode Outer DSCP Value is Copied to Decapsulated Packet DSCP Value for a Layer-3 Packet

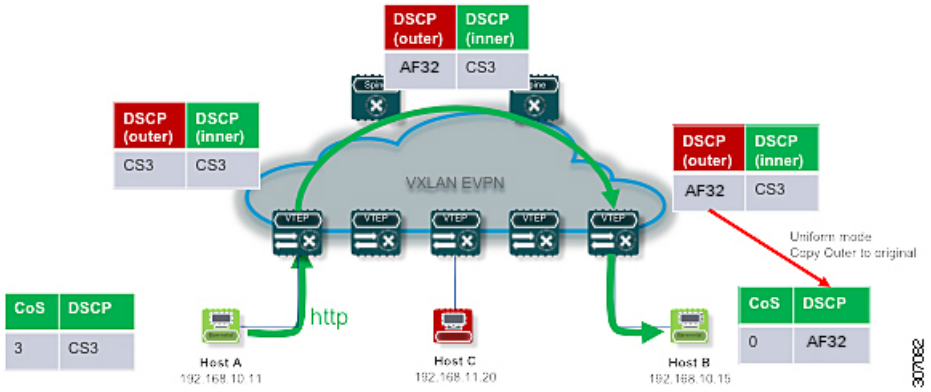

- Pipe mode – the original DSCP value is preserved at the VXLAN tunnel end. At the egress VTEP, the system copies the inner DSCP value to the decapsulated packet DSCP value. In this way, the original DSCP value is preserved at the end of the VXLAN tunnel.

Figure 6: Pipe Mode Inner DSCP Value is Copied to Decapsulated Packet DSCP Value for Layer-3 Packet

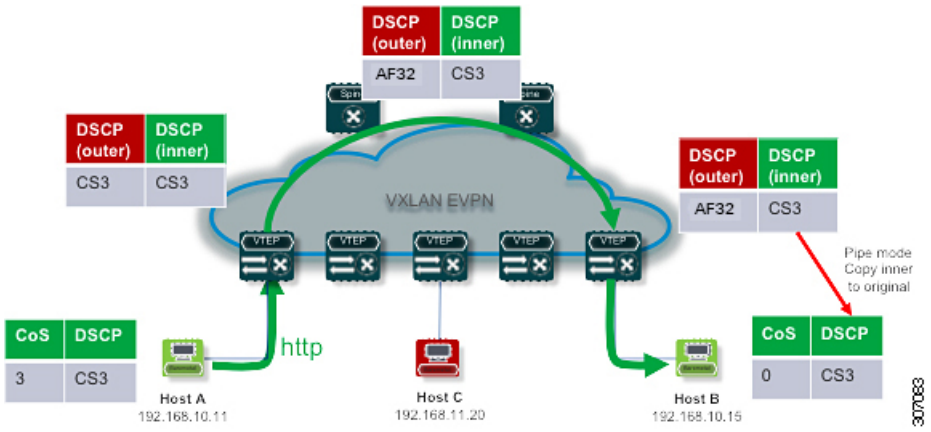

Guidelines and Limitations for VXLAN QoS

Platform and Release Support

- 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.

Unsupported Platforms and Line Card Exceptions

- 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 14.

Behavioral and Configuration Notes

- 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.

VXLAN QoS on the DCI Handoff Node

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.
- 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.

Outer DSCP Based VXLAN QoS Policy Feature
Outer DSCP Based VXLAN QoS Policy Guidelines and Limitations
VXLAN QoS Policies with Border Gateway (BGW) Spine

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.

Platform-Specific VXLAN QoS Support and Limitations

Default Settings for VXLAN QoS

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

Table 2: 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.

- Configuration of VXLAN QoS is done using the MQC model.
- The same configuration that is used for the QoS configuration applies to VXLAN QoS.
- Service policy must be attached to an NVE interface to account for all VXLAN traffic.

VXLAN QoS Configuration Reference

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

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.

a) **[no] match [access-group | cos | dscp | precedence]** *{name | 0-7 | 0-63 | 0-7}*

Example:

```
switch(config-cmap-qos)# match dscp 26
```

Configures traffic class by matching packets based on access-list, **cos**, **dscp**, or IP **precedence** values.

Step 3 **[no] policy-map type qos** *policy-map-name*

Example:

```
switch(config)# policy-map type qos policy
```

Creates or accesses the QoS policy map and enters policy-map mode.

a) **[no] class** *class-name*

Example:

```
switch(config-pmap-qos)# class class1
```

Creates a reference to the class and enters policy-map class configuration mode.

b) **[no] set qos-group** *qos-group-value*

Example:

```
switch(config-pmap-c-qos)# set qos-group 1
```

Sets the QoS group value. The value can range from 1 through 126.

c) **exit**

Example:

```
switch(config-pmap-c-qos)# exit
```

Exits class-map mode.

Step 4 **[no] interface nve** *nve-interface-number*

Example:

```
switch(config)# interface nve 1
```

Enters NVE interface configuration mode.

a) **[no] service-policy type qos input** *policy-map-name*

Example:

```
switch(config-if-nve)# service-policy type qos input policy
```

Applies the input service policy to the NVE interface.

b) (Optional) **[no] qos-mode [pipe]**

Example:

```
switch(config-if-nve)# qos-mode pipe
```

Enables pipe mode for decapsulated packet priority selection. The default is uniform mode.

Verify the VXLAN QoS Configuration

This task describes how to verify the VXLAN QoS configuration and view details of class maps, policy maps, and QoS settings on the switch.

SUMMARY STEPS

1. **show class-map**
2. **show policy-map**
3. **show running ipqos**

DETAILED STEPS

Procedure

Step 1 **show class-map**

Displays information about all configured class maps.

Step 2 **show policy-map**

Displays information about all configured policy maps.

Step 3 **show running ipqos**

Displays the configured QoS settings on the switch.

VXLAN QoS Configuration Examples

This topic provides configuration examples for VXLAN Quality of Service (QoS), including ingress and egress classification and marking, queuing, and CoS preservation on NVE interfaces.

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
```