



# Adaptive QoS

**Table 1: Feature History**

Feature Name	Release Information	Description
Adaptive QoS	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature enables WAN interface shapers and per-tunnel shapers at the enterprise edge to adapt to the available WAN bandwidth. The capability to adapt to the bandwidth controls differentiated packet drops at the enterprise edge and reduces or prevents packet drops in the network core.

- [Limitations and Restrictions, on page 1](#)
- [Information About Adaptive QoS, on page 2](#)
- [Configure Adaptive QoS, on page 4](#)
- [Configure Adaptive QoS Using the CLI, on page 5](#)
- [Customize Adaptive QoS Configuration, on page 6](#)
- [Monitor Adaptive QoS, on page 7](#)

## Limitations and Restrictions

- Adaptive QoS is only supported on Cisco IOS XE Catalyst SD-WAN devices. Only Cisco IOS XE Catalyst SD-WAN devices can be configured as hub devices. This means that for adaptive QoS to work, the spokes should also be Cisco IOS XE Catalyst SD-WAN devices.
- Adaptive QoS is only supported on to hub-to-spoke network topology.
- Adaptive QoS support on DIA/DCA interfaces is dependent upon the throughput of the overlay session.
- If an edge device is configured as hub, the WAN interface on the edge device cannot be configured with adaptive QoS.
- Adaptive QoS is supported for loopback interfaces only when a single loopback interface is bound to a single physical interface.

# Information About Adaptive QoS

## Overview of Adaptive QoS

Enterprise networks are increasingly using the Internet as a form of WAN transport. Therefore, QoS models need to adapt accordingly. QoS works effectively when deployed in a service-level agreement (SLA) environment, like Multiprotocol Label Switching (MPLS) networks. The available bandwidth on the Internet at a given time can vary. It can often be much lesser than the actual bandwidth that is offered by the service provider. In a non-SLA environment, QoS has limitations because it can't predict the changing bandwidth on the link.

With adaptive QoS, the shapers at the edge of the enterprise (WAN interface shaper and per-tunnel shaper) can adapt to the available WAN bandwidth, both Internet and Long-term Evolution (LTE). Thus, adaptive QoS can control differentiated drops at the enterprise edge and reduce the packet drops in the Internet core. When the adaptive QoS capability is not available, shapers that are applied as part of the egress QoS policy are static in value. They are configured based on the service provider bandwidth offering and don't change with time, thus they don't reflect the actual available Internet bandwidth.

### Benefits of Adaptive QoS

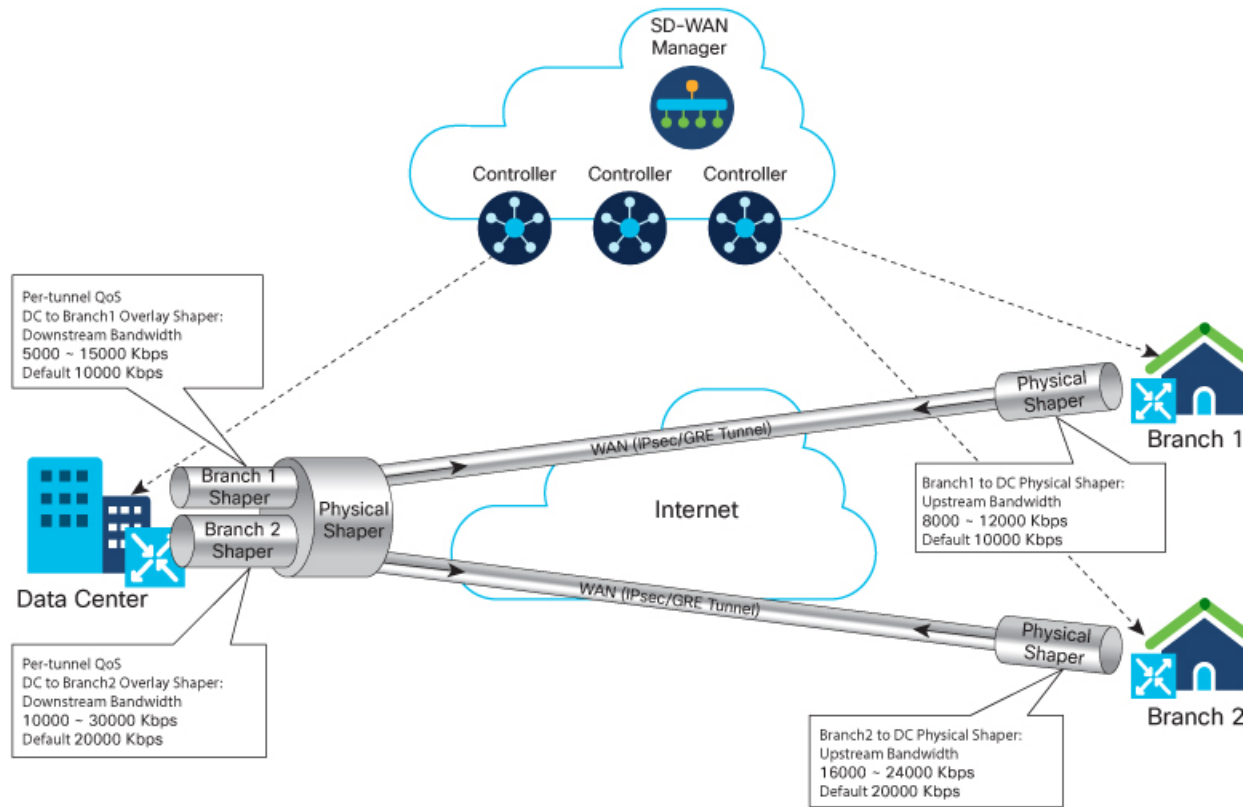
- Adjusts the shaper parameters based on the actual available Internet bandwidth in both directions, which is periodically computed
- Allows configuring a QoS policy on the spoke towards the hub
- Ensures better control of application performance at the enterprise edge even when the bandwidth fluctuates
- Allows aggregate tunnel shape adaptation to provide effective bandwidth between spoke and hub

## How Adaptive QoS Works in Cisco Catalyst SD-WAN

LTE and Internet bandwidth changes dynamically based on weather conditions and external parameters. In addition, Internet bandwidth can also fluctuate with the network conditions of the service provider, their congestion, and configurations.

Application traffic is prone to packet drops at the Internet core when the bandwidth is less, and the traffic can't be differentiated by user-defined priority. In such scenarios, Cisco Catalyst SD-WAN adaptive QoS automatically updates the shaper rate in real time based on the Internet and LTE link bandwidth.

Adaptive QoS can be enabled through Cisco SD-WAN Manager on a specific interface of an edge device that is configured with the spoke role in a hub-to-spoke network topology. You can specify minimum, maximum, and default values. You can also configure a timer interval at a global level on a WAN interface to measure the drop rates.



In the image, adaptive per-tunnel QoS is configured on the WAN interfaces of the spoke devices (Branch 1 and Branch 2) through Cisco SD-WAN Manager with the following configurations:

**Table 2: Branch 1 Configuration**

Parameter	Values
Upstream Bandwidth (Edge/spoke device in branch 2 to edge device/hub in the data center)	<ul style="list-style-type: none"> <li>• <b>Range:</b> 8000 Kbps—12000 Kbps</li> <li>• <b>Default:</b> 10000 Kbps</li> </ul>
Downstream Bandwidth (Edge device/hub of data center to the edge/spoke device in branch 2)	<ul style="list-style-type: none"> <li>• <b>Range:</b> 5000 Kbps—15000 Kbps</li> <li>• <b>Default:</b> 10000 Kbps</li> </ul>
Adapt Period	30 Minutes

Table 3: Branch 2 Configuration

Parameter	Values
Upstream Bandwidth (Edge/spoke device in branch 2 to edge device/hub in the data center)	<ul style="list-style-type: none"> <li>• <b>Range:</b> :16000 Kbps—20000 Kbps</li> <li>• <b>Default:</b> 10000 Kbps</li> </ul>
Downstream Bandwidth (Edge device/hub of data center to the edge/spoke device in branch 2)	<ul style="list-style-type: none"> <li>• <b>Range:</b> 10000 Kbps—30000 Kbps</li> <li>• <b>Default:</b> 20000 Kbps</li> </ul>
Adapt Period	60 Minutes

### Data Center

In the image, the per-tunnel QoS on the hub device in the data center reflects the downstream bandwidth configured on the spoke devices in Branch 1 and Branch 2.

Adaptive QoS in Cisco Catalyst SD-WAN is based on an algorithm based on packet drop or loss that works as follows:

Traffic Behavior	Analysis and Corresponding Adaptive QoS Behavior
No drops on the WAN or shapers	Shaper rate is not adjusted because the user traffic rate is low
Packet Drops in WAN	WAN bandwidth is oversubscribed. The shaper rate is thus adjusted to go DOWN to avoid packet drops.
Packet Drops in Shapers	A drop in shapers without packet drops in WAN indicates that the available WAN bandwidth is not being utilized; and therefore, the shaper rate is adjusted to go UP.

## Workflow of Adaptive QoS

When the adapt period is configured, adaptive QoS follows these stages.

- **Adapt:** This is the initial stage where the shaping rate is based on the default value or is recalculated based on the last cycle.
- **Measure:** In this stage, shaper or WAN loss metrics are calculated. A single adapt cycle can have multiple loss measurement cycles.
- **Verify or Recalculate:** This stage verifies whether the shaper rate works as expected. Based on both the shaper rate and WAN loss rate, the drop-based algorithm is used to calculate the appropriate shaping rate for the next cycle.

## Configure Adaptive QoS

To configure adaptive QoS use the Cisco VPN template for one of the following interfaces: Ethernet, Cellular, or DSL.

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
2. Click **Feature Templates** and then click **Add Template**.



---

**Note** In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is called **Feature**.

---

3. Choose a device from the list on the left. Feature templates that are applicable to the device are shown in the right pane.
4. Choose one of the available Cisco VPN Interface templates. In this example, we've chosen the **Cisco VPN Interface Ethernet** template.
5. Enter a name and description for the feature template.
6. Click **ACL/QoS**.
7. Notice that Adaptive QoS is disabled by default. To enable it, from the Adaptive QoS drop-down list, choose **Global**, and choose **On**.
8. (Optional) Enter adaptive QoS parameters. You can leave the additional details at as default or specify your values.
  - **Adapt Period:** Choose **Global** from the drop-down list, click **On**, and enter the period in minutes.
  - **Shaping Rate Upstream:** Choose **Global** from the drop-down list, click **On**, and enter the minimum, maximum, and default upstream bandwidth in Kbps.
  - **Shaping Rate Downstream:** Choose **Global** from the drop-down list, click **On**, and enter the minimum, maximum, downstream, and upstream bandwidth in Kbps.
9. Click **Save**.
10. [Attach the feature template to a device template.](#)

## Configure Adaptive QoS Using the CLI

The following example shows the adaptive QoS configuration on the Ethernet interface of a Cisco IOS XE Catalyst SD-WAN device.

```
sdwan
interface GigabitEthernet1
qos-adaptive
  period 90
  downstream 8000
  downstream range 6000 10000
  upstream 8000
  upstream range 4000 16000
exit
tunnel-interface
  encapsulation ipsec weight 1
  color biz-internet
  no last-resort-circuit
  vmanage-connection-preference 5
  allow-service all
```

```

no allow-service bgp
allow-service dhcp
allow-service dns
allow-service icmp
no allow-service sshd
no allow-service netconf
no allow-service ntp
no allow-service ospf
no allow-service stun
allow-service https
exit
exit

```

## Customize Adaptive QoS Configuration

### How the Adaptive QoS Algorithm Works

With adaptive QoS, the QoS shaper rate is adjusted based on WAN loss and local network loss. WAN and local network loss are measured based on the IPSec or GRE sequence number in the overlay session. The adaptive QoS algorithm brings the QoS shaper rate DOWN when WAN loss crosses the configured threshold. The algorithm brings the QoS shaper rate UP when the local network loss is detected to be over the threshold.

The adaptive QoS algorithm consists of the following parameters.

**Table 4:**

Parameter	Default Value	Customizable	Adaptive QoS Behavior
WAN loss threshold	0.2%	Yes	If the WAN loss detected is above the threshold, the shaper rate adapts DOWN.
Spoke overlay traffic threshold	40%	Yes	The QoS shaper rate for a spoke adapts UP or DOWN only when the overlay user traffic out of the overall traffic is above the threshold.
Local loss threshold	0.1%	No	If only local loss detected is above the threshold (WAN loss is still within the threshold), and the traffic throughput crosses a certain usage threshold, the shaper rate adapts UP.
Pre-adapting UP overlay user traffic threshold	90%	No	When both WAN and local loss are below the threshold, and the overlay user traffic usage of the QoS shaper is above the threshold, the QoS shaper rate adapts UP.
Hub adapting UP overlay user traffic usage threshold	90%	No	When only local loss is above the threshold, and the QoS shaper rate for and the overlay user traffic is above the threshold, the QoS shaper rate adapts UP. This behavior prevents adapting if there is congestion on the parent schedulers.

Parameter	Default Value	Customizable	Adaptive QoS Behavior
Adapting Stride	Minimum: 1% Maximum: current QoS shaper rate	No	The QoS shaper rate adapts UP if the normal stride is 1/10 of the QoS shaper rate range.  The QoS shaper rate adapts DOWN based on the normal stride in the shaper rate and WAN loss rate.

### Customize Adaptive QoS Thresholds

The following parameters in the adaptive QoS algorithm can be customized.

- **WAN loss threshold:**

Use the **platform qos sdwan adapt wan-loss-permillage** <1~999 permillage> through the configuration mode of the device CLI or using the [CLI add-on feature template](#) in Cisco SD-WAN Manager.

- **Spoke overlay traffic percentage:**

Use the **platform qos sdwan adapt spoke-overlay-usage** <1~100 percent> through the configuration mode of the device CLI or using the [CLI add-on feature template](#) in Cisco SD-WAN Manager.

## Monitor Adaptive QoS

### Verify Upstream Configuration

The following sample output shows the adaptive QoS statistics collected for upstream traffic.

```
Device# show platform software sdwan qos adapt stats
INTERFACE          DEFAULT  MIN      MAX      PERD    SHAPE-RATE
                   (kbps)  (kbps)  (kbps)  (min)   (kbps)
GigabitEthernet0/0/4 20000   10000   40000   1       40000
```

This sample output shows upstream adaptive QoS statistics.

```
Device# show platform software sdwan qos adapt history all

SDWAN upstream adaptive QoS
Interface: GigabitEthernet3
Adaptive QoS History:
TIME                LOCAL-LOSS  WAN-LOSS  TOTAL-OFFER  THROUGHPUT  ADAPT  SHAPE-RATE
                   (pps)      (kbps)   (pps)        (kbps)      (kbps)
2020-06-08T07:49:46 0.0%       0.0%     9600         13827       NOPE  50000
2020-06-08T07:48:46 0.0%       0.0%     9600         13826       NOPE  50000
2020-06-08T07:47:46 0.0%       0.0%     9600         13825       NOPE  50000
2020-06-08T07:46:46 0.0%       0.0%     9600         13827       NOPE  50000
2020-06-08T07:45:46 0.0%       0.0%     9600         13828       NOPE  50000
2020-06-08T07:44:46 0.0%       0.0%     9600         13828       NOPE  50000
2020-06-08T07:43:46 0.0%       0.0%     9600         13827       NOPE  50000
2020-06-08T07:42:46 0.0%       0.0%     9600         13832       NOPE  50000

SDWAN upstream adaptive QoS
Interface: Loopback0
Adaptive QoS History:
TIME                LOCAL-LOSS  WAN-LOSS  TOTAL-OFFER  THROUGHPUT  ADAPT  SHAPE-RATE
                   (pps)      (kbps)   (pps)        (kbps)      (kbps)
2020-06-08T07:49:46 50.8%      0.0%     16282        7980        UP    8099
```

```

2020-06-08T07:48:46 50.2%    0.8%    16282    8073    DOWN  8019
2020-06-08T07:47:46 50.8%    0.0%    16287    8005    UP    8099
2020-06-08T07:46:46 50.4%    0.7%    16282    8056    DOWN  8019
2020-06-08T07:45:46 50.9%    0.0%    16282    7976    UP    8099
2020-06-08T07:44:46 50.2%    0.9%    16282    8084    DOWN  8019
2020-06-08T07:43:46 50.7%    0.1%    16282    8002    UP    8099
2020-06-08T07:42:46 50.2%    0.9%    16282    8083    DOWN  8019

```

This sample output shows the history of the upstream adaptive QoS for the specified interface.

```

Device# show platform software sdwan qos adapt history GigabitEthernet0/0/4
SDWAN upstream adaptive QoS
Interface: GigabitEthernet0/0/4
Adaptive QoS History:
TIME                LOCAL-LOSS WAN-LOSS TOTAL-OFFER THROUGHPUT ADAPT SHAPE-RATE
                   (pps)      (kbps)      (kbps)
2020-05-21T02:43:44 56.0%    0.0%    34952    22087    UP    25100
2020-05-21T02:42:44 62.0%    0.0%    34952    19089    UP    22100
2020-05-21T02:41:44 67.9%    0.0%    34952    16091    UP    19100
2020-05-21T02:40:44 73.9%    0.0%    34952    13091    UP    16100
2020-05-21T02:39:44 79.9%    0.0%    34952    10091    UP    13100
2020-05-21T02:38:44 80.1%    0.0%    34952    9990     UP    10100
2020-05-21T02:37:44 80.1%    29.4%   34952    9990     DOWN 10000
2020-05-21T02:36:44 80.1%    29.4%   34952    9990     DOWN 10000

```

The following sample output shows the verification of the adaptive QoS configuration in the policy map applied to the GigabitEthernet3 interface.

```

Device# show policy-map interface GigabitEthernet3

Service-policy output: shape_GigabitEthernet3

Class-map: class-default (match-any)
 89140978 packets, 16580958431 bytes
 30 second offered rate 23246000 bps, drop rate 5255000 bps
Match: any
Queueing
queue limit 83 packets
(queue depth/total drops/no-buffer drops) 0/25186569/0
(pkts output/bytes output) 73636046/13696793305
shape (average) cir 20143000, bc 80572, be 80572
target shape rate 20143000

Service-policy : qos_policy_4class_cedge

queue stats for all priority classes:
  Queueing
  priority level 1
  queue limit 512 packets
  (queue depth/total drops/no-buffer drops) 0/0/0
  (pkts output/bytes output) 2004810/373383409

Class-map: Critical (match-any)
 6566981 packets, 1222194617 bytes
 30 second offered rate 13000 bps, drop rate 0000 bps
Match: qos-group 0
Priority: Strict, b/w exceed drops: 0

Priority Level: 1

```



## Verify Downstream Configuration

The following sample output shows the downstream adaptive QoS configuration.

```
Device# show sdwan omp tlocs

-----
tloc entries for 10.6.0.3
      lte
      ipsec
-----

                RECEIVED FROM:
peer           10.8.3.3
status         C,I,R
loss-reason    not set
lost-to-peer   not set
lost-to-path-id not set
Attributes:
  attribute-type  installed
  encap-key       not set
  encap-proto     0
  encap-spi       261
  encap-auth      sha1-hmac,ah-sha1-hmac
  encap-encrypt   aes256
  public-ip       176.16.6.2
  public-port     12346
  private-ip      176.16.6.2
  private-port    12346
  public-ip       176:16:6::2
  public-port     12346
  private-ip      176:16:6::2
  private-port    12346
  bfd-status      up
  domain-id       not set
  site-id         601
  overlay-id      not set
  preference      1000
  tag             not set
  stale           not set
  weight          1
  version         3
  gen-id          0x8000012c
  carrier         carrier6
  restrict        0
  on-demand       0
  groups          [ 0 ]
  bandwidth       80000
  bandwidth-dmin  50000
  bandwidth-down  100000
  bandwidth-dmax  100000
  adapt-qos-period 15
  adapt-qos-up    1
  qos-group       default-group
  border          not set
  unknown-attr-len not set
```

The following sample output shows downstream adaptive QoS statistics.

```
Device# show platform software sdwan qos adapt stats
===== Adaptive QoS Stats =====
COLOR          DEST-TLOC      DEST-IP        DEST-PORT  ENCAP  DEFAULT  MIN      MAX
PERD SHAPE-RATE                                     (kbps)   (kbps)   (kbps)
(min) (kbps)
```

```

lte          172.16.255.11  10.0.5.11      12347      IPSEC 100000    50000    100000
1           100000
lte          172.16.255.14  10.1.14.14     12346      IPSEC 100000    50000    100000
1           100000

```

The following sample output shows the adaptive QoS history of the eight most recent sessions on the tunnel.

```

Device# show platform software sdwan qos adapt history Tunnell 10.1.14.14
SDWAN OMP Session
Color: lte
Dest Tloc: 172.16.255.14
Dest IP: 10.1.14.14
Dest Port: 12346
Encap: IPSEC
Adaptive QoS History:
TIME                LOCAL-LOSS WAN-LOSS TOTAL-OFFER THROUGHPUT ADAPT SHAPE-RATE
                   (pps)          (kbps)
2020-05-21T04:51:28 30.0%      0.4%      87380      87852      DOWN 86973
2020-05-21T04:50:28 28.0%      2.9%      87380      90481      DOWN 87851
2020-05-21T04:49:28 31.9%      0.0%      87380      85553      UP   90474
2020-05-21T04:48:28 35.9%      0.0%      87380      80477      UP   85474
2020-05-21T04:47:28 39.9%      0.0%      87380      75475      UP   80474
2020-05-21T04:46:28 40.5%      0.0%      87380      74727      UP   75474
2020-05-21T04:45:28 39.9%      0.6%      87380      75480      DOWN 74727
2020-05-21T04:44:28 40.5%      0.0%      87380      74737      UP   75481

```

The following sample output shows the average shaper rate, target shaper rate, and the bandwidth remaining.

```

Device# show policy-map multipoint Tunnell 10.1.14.14

Interface Tunnell <--> 10.1.14.14

Service-policy output: SDWANPolicy4325397

Class-map: class-default (match-any)
 343994858 packets, 59167000300 bytes
 5 minute offered rate 114034000 bps, drop rate 37596000 bps
Match: any
Queueing
queue limit 362 packets
(queue depth/total drops/no-buffer drops) 0/76866318/0
(pkts output/bytes output) 246459053/45841211875
shape (average) cir 87966000, bc 351864, be 351864
target shape rate 87966000
bandwidth remaining ratio 8

Service-policy : qos_policy_4class_cedge

queue stats for all priority classes:
Queueing
priority level 1
queue limit 512 packets
(queue depth/total drops/no-buffer drops) 0/0/0
(pkts output/bytes output) 40145494/7466889901

Class-map: Critical (match-any)
 68807464 packets, 11834768360 bytes
 5 minute offered rate 22815000 bps, drop rate 10139000 bps
Match: qos-group 0
Priority: Strict, b/w exceed drops: 0

Priority Level: 1
police:
  rate 15 %

```

```
rate 13065500 bps, burst 408296 bytes  
conformed 40141805 packets, 6904295187 bytes; actions: transmit
```

