

SD-WAN Control Traffic Overhead User Guide

Contents

[Introduction](#)

[Problem](#)

[Solution](#)

[Generic Guideline for Overhead Calculation](#)

[Example for Overhead Calculation](#)

Introduction

This document describes how to calculate control traffic overhead on an SD-WAN overlay deployment. Please note the following article guidance should be used on viptela code below 20.10.x and IOS-XE SD-WAN 17.10.x and below (from 20.10.x /17.10.x Cisco has implemented push model for data collection).


Problem

A common question that is received at the time of the design phase from a user is 'How much overhead the SD-WAN solution would make to our branch circuit'? The answer is that it depends on a few variables.

Solution


This case study helps you find that answer. Most of the users at the time of a branch roll out can or can not have the internet circuit provisioned. If they have one it would typically look something like Figure 1.

Figure1. SD-WAN Branch with both internet and Multiprotocol Label Switching (MPLS) circuit.




This might not be the case always, some users would mostly prefer to migrate to SD-WAN with minimal change and new circuit introduction, the addition of circuit possibly planned for a later phase, which would be like Figure 2. without an Internet circuit.

Figure 2. SD-WAN Branch with only MPLS circuit.



In order to set the stage, if you have 100 branches with 2 Head Ends, and a proposed full-mesh topology between branches and Head Ends, and the user has a strict QOS standard with 20% allocation to Low Latency Queue (LLQ) for voice.

With the migration to SD-WAN what would be our overhead to consider for these branches, if so any. Let's dig deeper.

 **Note:** These calculations are to be considered at a normal operational requirement inclusive of peak requirement. However, don't consider all possible scenarios.

These numbers are derived from the lab test that was performed with 1vManage, 1vBond, and 1vSmart, 255 BFD sessions.

Table 1. Bandwidth per session.

| | |
|------------------------|--|
| 1 BFD Session/Neighbor | $2 \times 132 \times 8 = 2.2 \text{ Kbps}$ 2: In a second you send and receive up to 2 BFD packets 132: BFD packet size in B |
|------------------------|--|

| | |
|--------------------------|------------------|
| DTLS to vSmart | up to 80 Kbps* |
| vManage polling for data | up to 1.2 Mbps** |
| Enabling DPI | 200 Kbps |

Kbps = Kilobits per second

B = Bytes

Mbps = Megabits per second

* Depends on the policy and routes; this calculation is needed only at the time of the initial exchange and the stable state is much lower/minimal around 200 B.

** Doesn't consider a user-triggered activity such as running remote commands or admin tech; 1.2 Mbps is at peak spike.

Now, if you consider all the 100 full mesh sites which are 200 BFD sessions (2 routers per branch, 2 TLOC's per router with **restrict** in color), the previously mentioned table would become.x.

Table 2. Queue0 Bandwidth for 200 BFD sessions [100 Sites] that includes vSmart and vManage polling.

| | |
|-----------------|-------------------------|
| 200 BFD Session | 440 Kbps [2.2 x 200] |
| DTLS to vSmart | up to 80 Kbps* |
| vManage polls | up to 1.2 Mbps** |
| Total | 1.72 Mbps |

* Depends on the policy and routes; this calculation is needed only at the time of the initial exchange and the stable state is much lower/minimal around 200 B.

** Doesn't consider a user-triggered activity such as running remote commands or admin tech; 1.2 Mbps is at peak spike.

Keep this in mind all these traffic hits Queue0 LLQ, this control traffics are always given first-class citizen priority meaning they are the last one to be policed on an LLQ.

Often at the time of the QoS design, voice traffic is placed into Queue0 (LLQ), with a 1.72 Mbps requirement for 100 branches full mesh with Tloc for SD-WAN, you can see policing/drop on LLQ with low bandwidth circuit branches.

Now, if you consider the Tloc extension overhead which won't contribute towards Queue0 but constitutes the overall capacity requirement.

Table 3. Overall Bandwidth requirement after you consider how to control traffic over the Tloc extension.

| | |
|--|--------------------------------------|
| Queue0 Requirement | 1.72 Mbps |
| 200 BFD Session for Tloc Extension [Encrypted] Non Queue0 | 520 Kbps [440 + 80*] [BFD + DTLS] |
| Total | 2.24 Mbps |

* Depends on the policy and routes; this calculation is needed only at the time of the initial exchange and the stable state is much lower/minimal around 200 B.

Per 100 branch full meshed with TLOC extensions with color restrict consider a capacity planning of ~2.5 Mbps on an extreme requirement, again you can collect realtime commands, admin tech is not considered in the previously mentioned calculation, consider this in a normal operation situation.

Scenario 1.

If you need to accommodate control traffic requirements to Queue0 and if a branch has only a 10 Mbps circuit, that needs to be onboarded into SD-WAN overlay with a QoS policy of only 20% LLQ for both Voice and Control traffic. You can look at a degraded experience at the time of the peak polling from vManage. A hub and spoke solution might not help in this case since it still consumes around 1.28 Mbps.

Table 4. Hub and Spoke Queue0 Bandwidth requirement.

| | |
|----------------------------|-----------------------|
| 4 BFD Session to Head Ends | 8.8 Kbps [2.2 x 4] |
| DTLS to vSmart | upto 80 Kbps* |
| vManage polls | upto 1.2 Mbps** |
| Total | 1.28 Mbps |

* Depends on the policy and routes; this calculation is needed only at the time of the initial exchange and the stable state is much lower/minimal around 200 B.

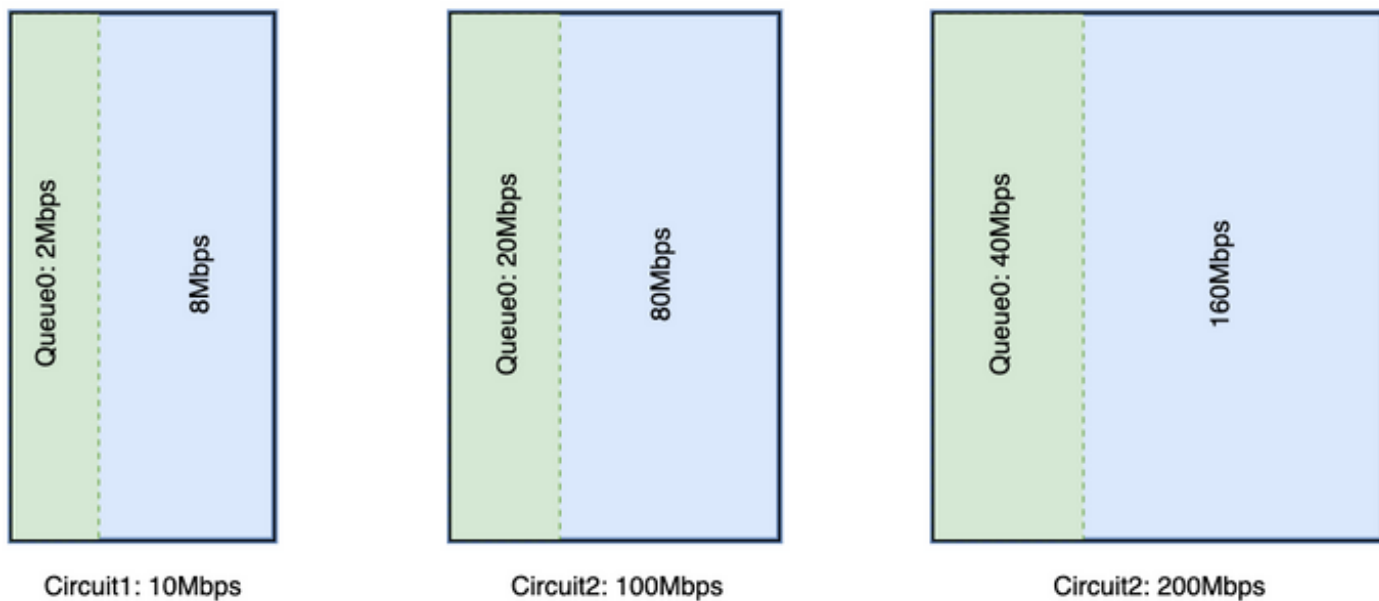
** Doesn't consider a user-triggered activity such as running remote commands or admin tech; 1.2 Mbps is at peak spike.

Scenario 2.

If you decide to redesign the QoS policy, to accommodate the ~2Mbps extra bandwidth requirement, you can increase the QoS LLQ from 20% to 40%. However, that would have a negative effect on larger bandwidth circuits.

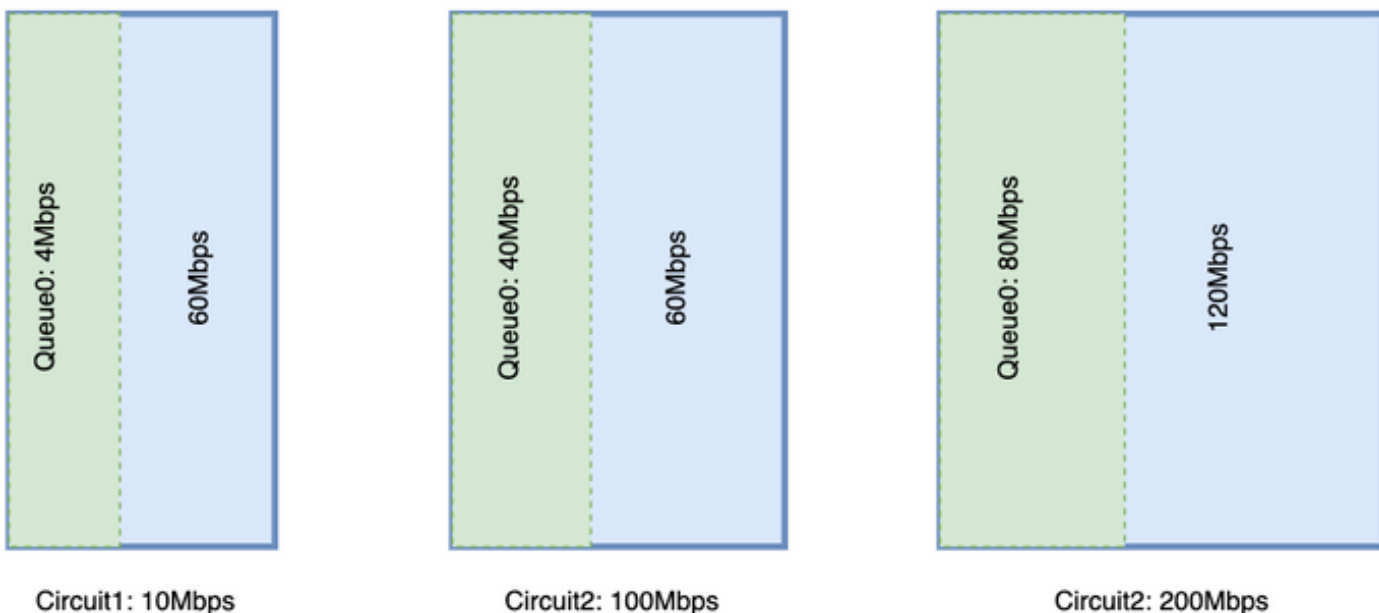
Figure 3. Typical 20% Queue0 Allocation for QoS.

Queue 0 at 20%



For a 10 Mbps circuit, Queue0 gets 2 Mbps at 20%. Assume this is a typical QoS standard for a company. SD-WAN adoption requires full-mesh, so, you need to increase the allocation of Queue0 to accommodate 2 Mbps overhead to Queue0 if the user decides to increase the QoS allocation to 40% as shown in the image.

Queue 0 at 40%




See that an enormous amount of Queue0 for a circuit takes away the resources for the other queue. However, the difference is more on a larger bandwidth circuit.

You must ideally have the LLQ to have a fixed allocation for the control traffic and another queue for voice traffic, but both require a priority queue. Cisco routers do support a priority queue with two levels known as split LLQ, again this doesn't address a minimum bandwidth requirement issue once a minimum requirement is met a split LLQ would be a preferred QoS design

Split LLQ:

Queue 0 at 20%



With Split LLQ, you add the necessary bandwidth to the Queue and still maintain the priority queue.

The split LLQ currently supports only with addon CLI, with split LLQ could have two levels of the priority queue, a sample configuration would be as shown here. The configuration can be customized with variables, this snippet reserves 4 Mbps for the control traffic and the rest of the queue as assigned bandwidth percentage.

Example for a Split Queue:

```
<#root>
policy-map GBL_edges_qosmap_rev1
class Queue0

priority level 1

    police cir 2000000 bc 250000
    conform-action transmit
    exceed-action drop
!
```

```
!  
class Queue1  
    bandwidth remaining ratio 16  
    random-detect precedence-based  
!  
class class-default  
    bandwidth remaining ratio 8  
    random-detect precedence-based  
!  
class Queue3  
    bandwidth remaining ratio 16  
    random-detect precedence-based  
!  
class Queue4  
    bandwidth remaining ratio 32  
    random-detect precedence-based  
!  
class Queue5  
    bandwidth remaining ratio 8  
    random-detect precedence-based  
!  
class Queue6  
  
priority level 2  
  
    police rate percent 20  
!  
!  
!
```

 **Note:** These configurations are tested on ISR/ASR running 17.3.x and Controllers on 20.3.x.

Generic Guideline for Overhead Calculation

This table can help you plan capacity per circuit for an SD-WAN control overhead.

Table 5. Generic guideline calculation (assumes you have color restrict).

| Protocol/Session | Bandwidth Required |
|---------------------------|--|
| Queue0 | $2.2 \times [\text{no.of Sites} \times \text{no.of BFD to a site from WAN Tloc}] + 80 + 1200$ $\text{BFD Size} \times [\text{no.of Sites} \times \text{no.of BFD to a site from WAN Tloc}] + \text{DTLS} + \text{vManage}$ $= \text{Queue0_Allocation}$ |
| Control traffic over TLOC | $2.2 \times [\text{no.of Sites} \times \text{Tloc/per router}] + 80$ $\text{BFD Size} \times [\text{Sites} \times \text{TLOC/per router}] + \text{DTLS}$ $= \text{Tloc_Allocation}$ |
| Total | $\text{Queue0_Allocation} + \text{Tloc_Allocation}$ |

Example for Overhead Calculation

If you need to calculate the overhead of the MPLS circuit for 100 sites similar to the one shown here, you can assume that each color has restrict enabled.

No. of Sites = 100

No. of BFD to a site from WAN Tloc = 2.

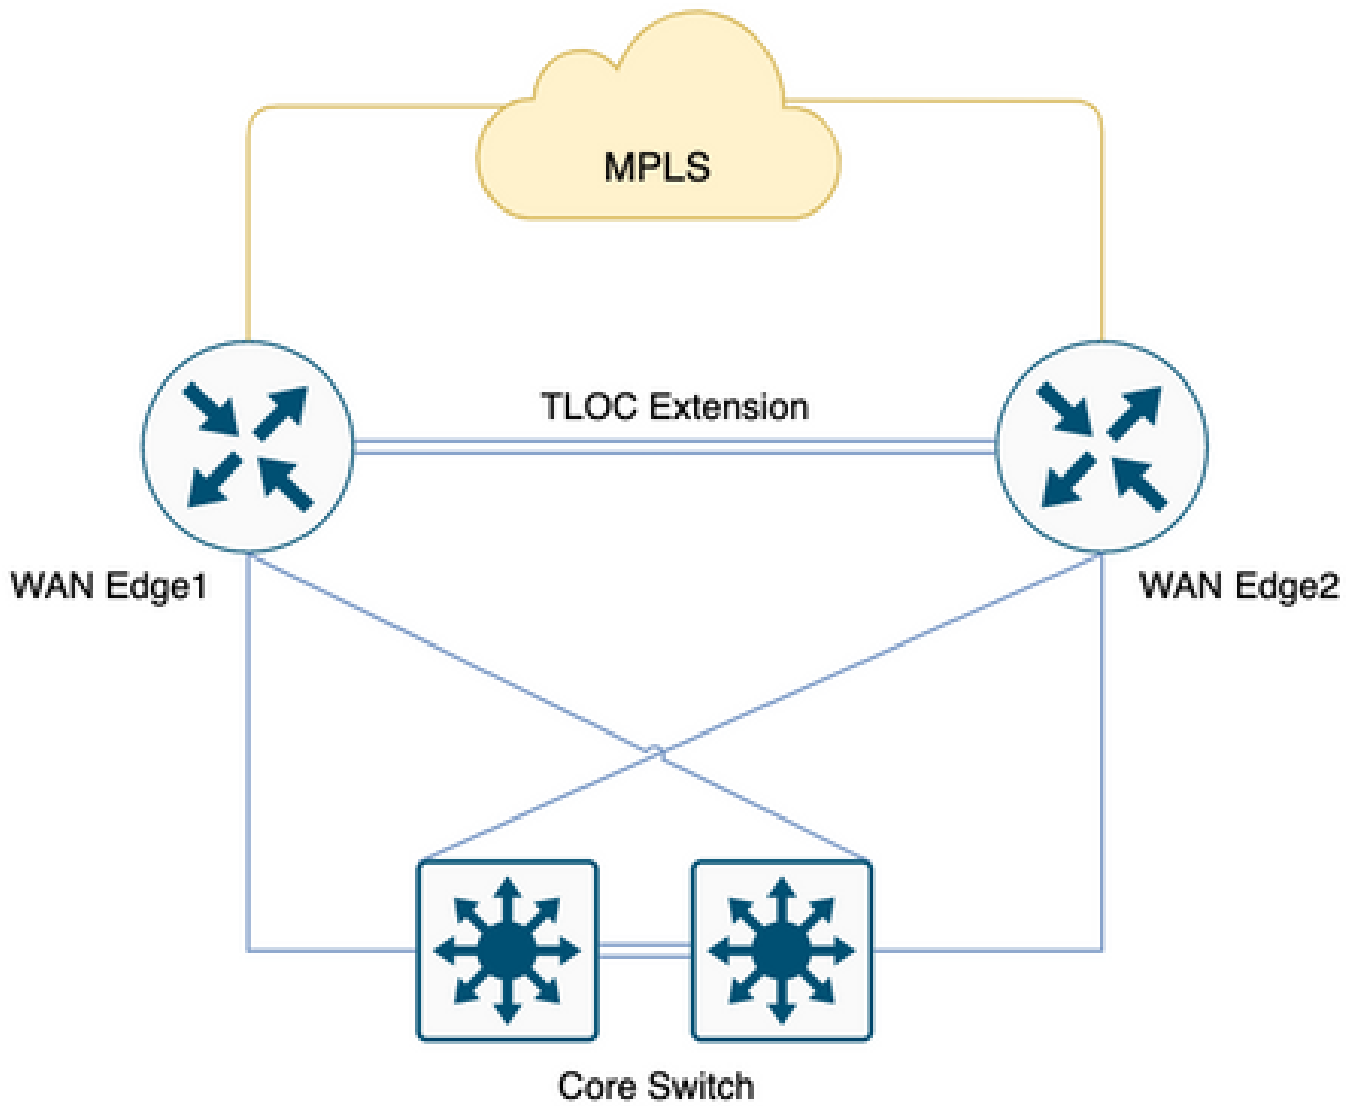


Table 6. Calculate the MPLS overhead for the deployment of 100 Sites.

| Protocol/Session | Bandwidth Required |
|---------------------------|---|
| Queue0 | $2.2 \times [100 \times 2] + 80 + 1200$ BFD Size x [no.of Sites x no.of BFD to a site from WAN Tloc] + DTLS +vManage = 1.72 Mbps |
| Control traffic over TLOC | $2.2 \times [100 \times 2] + 80$ BFD Size x [Sites x TLOC/per router] + DTLS = 520 Kbps |
| Total | 1720 Kbps + 520 Kbps |

| | |
|--|-------------|
| | = 2.24 Mbps |
|--|-------------|

Queue0 overhead of 1.72 Mbps and total overhead is 2.24 Mbps.