

Understand BFD Protocol Relationship with App-Aware Routing

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Tunnel Performance Statistics Calculation](#)

[Examples of BFD Values Relation with Loss](#)

Introduction

This document describes the relationship that exists between the BFD Hello packets and the App-Aware Routing Tunnel statistics.

Prerequisites

Requirements

Cisco recommends that you have knowledge of this topics:

- Cisco Catalyst Software-Defined Wide Area Network (SD-WAN).
- App-Aware Routing.
- BFD.

Components Used

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

- Cisco Catalyst SD-WAN Manager.
- Cisco IOS® XE Catalyst SD-WAN Edges.

Background Information

The Bidirectional Forwarding Detection (BFD) protocol runs over all data plane tunnels between Cisco IOS-XE Catalyst SD-WAN devices. This protocol is used to monitor the liveness and path characteristics of the tunnels such as the tunnel performance reported as Loss, Jitter, and Latency.

Edge devices use BFD Hello probes to provide a measurement of packet loss, jitter, and latency on the tunnel. These statistics are calculated for each BFD Hello probe and are taken over a sliding window of time

called polling interval.

These loss, latency, and jitter statistics are used by App-Aware Routing to deliver the traffic based on requirements set in the policy, called SLA classes, in which it determines the maximum loss, jitter, and latency allowed in the tunnel selected to deliver the data.

Because of this, it is highly important to understand how the measures are calculated and how a change in the BFD values can affect the tunnel performance calculation principally mean loss. The BFD parameters are:

Parameter	Default Value	Range	Use
BFD hello interval	1 second	1 through 65535 seconds	Packets to detect the liveness of the tunnel connection and to detect faults on the tunnel.
Polling interval	10 minutes (600,000 milliseconds)	1 through 4,294,967 milliseconds	How often a bucket measure is calculated to provide statistics.
Multiplier	6	1 through 6	Value that multiplies polling interval to specify the time to calculate mean loss, mean latency, and mean jitter. This value determines the number of buckets.

Tunnel Performance Statistics Calculation

For BFD parameters set as default, the calculation of the statistics is done as follows:

Polling Interval / BFD Hello Interval = 600,000 ms / 1000 ms = 600 BFD Hellos per bucket.

As the multiplier is set as 6, it means that 6 buckets are used to calculate the average latency, jitter, and loss. With default values this equals **1 hour**. This total time is also known as the app-route interval.

App-route interval = Polling interval * multiplier = 600,000 ms x 6 = 3,600,000 ms equal 1 hour.

The calculations of the App-route statistics are used by App-Aware Routing to determine changes in the data plane. In order for an Edge device to take advantage of the app-route statistics, SLA classes must be specified in the AAR policy in which maximum acceptable packet jitter, loss, and latency is set. These SLA classes are used in the AAR policy to route traffic for specified applications according to the SLAs.

Once configured in an Edge device, AAR statistics are used to compare against mean loss, mean latency, and mean jitter provided by the statistics calculated with all buckets (over the entire app-route interval). It is also important to note that the SLAs are updated after each poll-interval, every ten minutes by default.

To obtain mean loss, mean jitter, and mean latency the equations used are:

Mean Loss = (total loss across all buckets * 100) / Total packets.

Mean Latency = (total loss across all buckets) / amount of buckets.

Mean Jitter = (total jitter across all buckets) / amount of bucket.

The calculation of those values alongside the average of each bucket can be reviewed in the CLI with:

```
<#root>
```

```
vEdge#
```

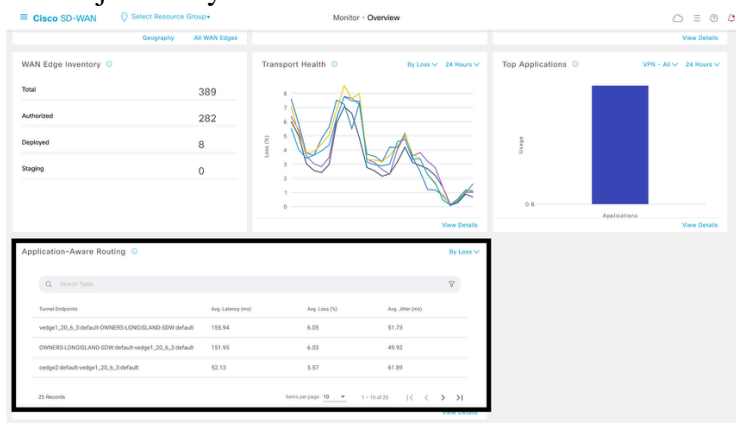
```
show app-route stats
```

```
<#root>
```

```
cEdge#
```


```
show sdwan app-route stats
```

While in the GUI, the mean loss, mean latency, and mean jitter only can be reviewed in **Monitor >**



Overview > Application-Aware Routing section.

It can also be reviewed in the **Monitor > Devices > Select Device > WAN > Tunnel** section.



Examples of BFD Values Relation with Loss

As BFD Hellos are configurable values, they can be modified based on requirements; however, it is important to modify them after careful consideration, otherwise skewed calculations or false positive statistics can be received since the accuracy of mean loss calculation depends on the BFD values. For example, with default values of:

Parameter	default
BFD hello packet	1 second
Polling interval	(600,000 milliseconds) 10 minutes
Multiplier	6

```
<#root>
```

```
vEdge1#
```

```
show app-route stats
```

```
app-route statistics 10.100.100.2 10.200.200.4 ipsec 12366 12346
remote-system-ip 10.1.1.1
```

```

local-color    private1
remote-color   private1
mean-loss      1
mean-latency   110
mean-jitter    51
sla-class-index 0,2
IPV6 TX  IPV6 RX

```

TOTAL INDEX	PACKETS	AVERAGE LOSS	AVERAGE LATENCY	TX DATA JITTER	RX DATA PKTS	DATA PKTS	DATA PKTS	DATA PKTS
0	596	7	110	50	0	0	0	0
1	596	5	111	50	0	1	0	0
2	597	13	111	53	0	0	0	0
3	594	4	111	53	0	0	0	0
4	596	5	110	50	0	0	0	0
5	594	12	111	50	0	2	0	0

$$\text{Mean Loss} = ((7+5+13+4+5+12)100) / (596+596+597+594+596+594)$$

$$= 4600/3573$$


$$= \mathbf{1.28 \sim 1\%}$$

$$\text{Mean Latency} = (110+111+111+111+110+111)/6$$

$$= \mathbf{110.66 \sim 110 \text{ ms}}$$

$$\text{Mean Jitter} = (50+50+53+53+50+50) / 6$$

$$= \mathbf{3 / 6 = 51 \text{ ms}}$$

 **Note:** For each calculation done, only integer values are presented. Even when decimal is the exact result, integer values are rounded to the lower nearest integer.

Normally, is a good option to modify these values to make the calculation more often but it can cause significant impact; for example, if instead of default values, the polling interval is modified to:

Parameter	default
BFD hello packet	1 second
Polling interval	(60,000 milliseconds) 1 min
Multiplier	6

This change means that it uses $1 \times 60 = 60$ packets per bucket instead of 600 as default. The result of mean loss is:

<#root>

vEdge1#

show app-route stats

app-route statistics 10.100.100.2 10.200.200.4 ipsec 12366 12346

remote-system-ip 10.1.1.1

local-color private1

remote-color private1

mean-loss 3

mean-latency 112

mean-jitter 51

sla-class-index 0,2

IPV6 TX IPV6 RX

TOTAL INDEX	PACKETS	AVERAGE LOSS	AVERAGE LATENCY	TX DATA JITTER	RX DATA PKTS	DATA PKTS	DATA PKTS	DATA PKTS
0	59	1	113	53	0	0	0	0
1	60	3	111	52	0	1	0	0
2	59	1	111	51	0	1	0	0
3	60	3	111	50	0	1	0	0
4	60	2	115	50	0	0	0	0
5	59	1	111	50	0	2	0	0

$$\text{Mean Loss} = ((1+3+1+3+2+1)*100)/(59+60+59+60+60+59)$$

$$= (1100)/ 357$$

$$= 3.08 \sim 3\%$$

At this point if for example, the SLA class is set a Maximum Loss of 3, then the tunnel is under the limit of the violation of the SLA. However, if the polling interval is modified to:

Parameter	default
BFD hello packet	1 second
Polling interval	(6,000 milliseconds) 1 second
Multiplier	6

This change means that it uses $1 \times 6 = 6$ packets per bucket instead of 600 as default. The result of mean loss is:

<#root>

vEdge1#

```
show app-route stats
```

```
app-route statistics 10.100.100.2 10.200.200.4 ipsec 12366 12346
```

```
remote-system-ip 10.1.1.1
```

```
local-color private1
```

```
remote-color private1
```

```
mean-loss 17
```

```
mean-latency 110
```

```
mean-jitter 0
```

```
sla-class-index None
```

```
IPV6 TX IPV6 RX
```

TOTAL INDEX	PACKETS	AVERAGE LOSS	AVERAGE LATENCY	TX DATA JITTER	RX DATA PKTS	DATA PKTS	DATA PKTS	DATA PKTS
-------------	---------	--------------	-----------------	----------------	--------------	-----------	-----------	-----------

0	5	1	113	2	0	0	0	0
---	---	---	-----	---	---	---	---	---

1	6	1	110	1	0	1	0	0
---	---	---	-----	---	---	---	---	---

2	6	1	111	2	0	0	0	0
---	---	---	-----	---	---	---	---	---

3	6	0	111	0	0	0	0	0
---	---	---	-----	---	---	---	---	---

4	6	1	111	0	0	0	0	0
---	---	---	-----	---	---	---	---	---

5	6	1	111	0	0	2	0	0
---	---	---	-----	---	---	---	---	---

Mean Loss = ((5)100)/(5+6+6+6+6+6)

= (500)/29

= **17.24 ~ 17%**

Whether the the poll interval is reduced without the correct validation of how many packets are used to measure, it can affect the mean loss, the same can apply if the bfd hello-interval is increase without increase the pool interval.

In the last example, as very few packets are used to make the calculation, with only one packet lost, the mean loss can be affected significantly. The result of those calculations is an App-Aware policy behavior with multiple and very frequent fail overs.

The purpose of this explanation is not to avoid the modification of those values, on the contrary, in many situations those probes are needed to be modified. This depends completely on network requirments, but it is highly important to review how much those hello packets can be decreased.

The configuration command to modified globally the poll interval is:

```
<#root>
```

```
vEdge(config)#
```

```
bfd app-route poll-interval 600000
```