



MEF 105 Draft (R4)

Performance Monitoring and Service Readiness Testing for SD-WAN

September 2022

**This draft represents MEF work in progress and is
subject to change.**

This draft document represents MEF work in progress; it has not achieved full MEF standardization and is subject to change. Changes are likely before this becomes a fully endorsed MEF Standard. The reader is strongly encouraged to keep this in mind and review the Release Notes (if applicable) when making a decision on adoption. Additionally, because this document has not been adopted as a Final Specification in accordance with MEF's Bylaws, Members are not obligated to license patent claims that are essential to implementation of this document under MEF's Bylaws.

Disclaimer

© MEF Forum 2022. All Rights Reserved.

The information in this publication is freely available for reproduction and use by any recipient and is believed to be accurate as of its publication date. Such information is subject to change without notice and MEF Forum (MEF) is not responsible for any errors. MEF does not assume responsibility to update or correct any information in this publication. No representation or warranty, expressed or implied, is made by MEF concerning the completeness, accuracy, or applicability of any information contained herein and no liability of any kind shall be assumed by MEF as a result of reliance upon such information.

The information contained herein is intended to be used without modification by the recipient or user of this document. MEF is not responsible or liable for any modifications to this document made by any other party.

The receipt or any use of this document or its contents does not in any way create, by implication or otherwise:

- a) any express or implied license or right to or under any patent, copyright, trademark, or trade secret rights held or claimed by any MEF member which are or may be associated with the ideas, techniques, concepts, or expressions contained herein; nor
- b) any warranty or representation that any MEF members will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor
- c) any form of relationship between any MEF member and the recipient or user of this document.

Implementation or use of specific MEF standards, specifications, or recommendations will be voluntary, and no Member shall be obliged to implement them by virtue of participation in MEF Forum. MEF is a non-profit international organization to enable the development and worldwide adoption of agile, assured, and orchestrated network services. MEF does not, expressly, or otherwise, endorse or promote any specific products or services.

Table of Contents

1	List of Contributing Members	1
2	Abstract.....	1
3	Terminology and Abbreviations.....	2
4	Compliance Levels	6
5	Conventions	7
5.1	Numerical Conventions	7
5.2	Diagram Conventions	7
6	Introduction.....	0
7	SD-WAN Performance Monitoring.....	2
7.1	Performance Monitoring Framework	2
7.2	Performance Monitoring Metrics	3
7.2.1	Application Monitoring Flow and TVC PM Metrics	3
7.2.2	PM Configuration	15
7.2.2.1	<i>Monitored Entity</i>	15
7.2.2.2	<i>PM Metric List</i>	15
7.2.2.2.1	PM Metrics for Application Monitoring Flows	16
7.2.2.2.2	PM Metrics for Application Flows	17
7.2.2.2.3	PM Metrics for TVCs	18
7.2.2.3	<i>PM Metric Calculation Duration</i>	19
7.2.2.4	<i>PM Metric Calculation Instance</i>	20
7.3	PM Metric Calculation and Reporting.....	22
7.3.1	PM Metric Calculation Data Set	25
7.3.2	PM Metric Report	26
7.4	Threshold Crossing Alerts	28
7.4.1	Stateful Threshold Crossing Alert Reporting	29
7.4.2	Stateless Threshold Crossing Alert Reporting.....	32
7.4.3	Threshold Crossing Alert Requirements.....	36
7.4.3.1	<i>Common TCA Requirements</i>	36
7.4.3.2	<i>Stateful TCA Requirements</i>	37
7.4.3.3	<i>Stateless TCA Reporting Requirements</i>	40
8	Service Readiness Testing for an SD-WAN Service	44
8.1	Verifying ordered pair of UCS End Points Connectivity	44
8.2	Service Readiness Testing Terms and Components	45
8.3	Service Readiness Measurement Point Locations	46
8.4	UCS Service Attribute Reporting	47
8.5	Service Readiness Testing Parameters and Results.....	47
8.5.1	SD-WAN Edge to SD-WAN Controller/Orchestrator Communication	47
8.5.2	Ordered Pairs of UCS End Points SRT.....	48
8.6	Test Methodology	48
8.6.1	UCS End Point Ordered Pair Connectivity Testing.....	49
8.7	Test Record.....	50
9	References.....	51



Appendix A	Service Readiness Testing vs Service Activation Testing	52
-------------------	--	-----------

List of Figures

Figure 1 – Diagram Conventions	0
Figure 2 – AFMP Locations	4
Figure 3 – Measured Information Rate Byte Counter Locations	5
Figure 4 – Ingress AFMP to Egress AFMP ordered pair Example	5
Figure 5 – TVC MP Locations	13
Figure 6 – Ingress TVC MP to Egress TVC MP ordered pair Example	14
Figure 7 – PM Metric Calculation and Reporting Example	23
Figure 8 – PM Metric Calculation Data Set Duration Equals the PM Metric Report Duration Example	24
Figure 9 – PM Metric Report Interval Greater Than PM Metric Calculation Data Set Duration	25
Figure 10 – Stateful TCA Flowchart	30
Figure 11 – Stateful TCA	31
Figure 12 – Stateful TCA without Consecutive PM Metric Calculation Intervals \geq TCA Performance Threshold Value Example	32
Figure 13 – Stateless TCA w/o Damping Factor Flowchart	33
Figure 14 – Stateless TCA w/Damping Factor Flowchart	34
Figure 15 – Stateless TCA Example	35
Figure 16 – Stateless TCA with Damping Factor Example	35
Figure 17 – SRT Order Pairs of UCS End Points Connectivity Verification Example	45
Figure 18 – SRMP Location of UCS End Points Connectivity Testing	47
Figure 19 – UCS End Point Ordered Pair Single Direction Connectivity SRT Methodology	48

List of Tables

Table 1 – AF Specification, zoneName and PM Metric Agreement Example	6
Table 2 – Monitored SD-WAN UNI ordered pair Agreement Example	7
Table 3 – Monitored SD-WAN UNI Agreement Example	7
Table 4 – AF Specification, zoneName, ordered pair, and PM Metric Example	9
Table 5 – AF Specification zoneName, UNI Identifier, and PM Metric Overview Example	11
Table 6 – PM Metric Calculation Instance Attributes	22
Table 7 – PM Metric Calculation Data Set Attributes	26
Table 8 – PM Metric Report Attributes	27
Table 9 – Stateful TCA Reporting PM Metric Parameters	38
Table 10 – Stateful TCA Reporting SET & CLEAR Criteria	38
Table 11 – Stateful TCA Notification Message Fields	40
Table 12 – Stateless TCA Reporting PM Metric Conditions	40
Table 13 – Stateless TCA Reporting Notification Message Fields	42
Table 14 – Damping Factor TCA Notification Message Field	43
Table 15 – UCS End Point Ordered Pair Connectivity Test Methodology	50

1 List of Contributing Members

The following members of the MEF participated in the development of this document and have requested to be included in this list.

Editor Note 1: This list will be finalized before Letter Ballot. Any member that comments in at least one CfC is eligible to be included by opting in before the Letter Ballot is initiated. Note it is the MEF member that is listed here (typically a company or organization), not their individual representatives.

- ABC Networks
- XYZ Communications

2 Abstract

An SD-WAN Service, as defined in MEF 70.1 [13], can use Performance Monitoring metrics to identify degradations or failures of TVC and Application Flows. The Performance metrics that are shared with SD-WAN Subscribers are defined in MEF 70.1 [13]. This document defines how the Performance Monitoring Metrics are measured and calculated. SD-WAN Performance Monitoring uses IP Packets to make performance measurements and uses these measurements to calculate Performance Monitoring Metric Values that can be reported.

The difference between the Performance metrics defined in MEF 70.1 [13] and the Performance Monitoring metrics defined within this document involve where the metrics are specified from. The MEF 70.1 [13] Performance metrics are defined from SD-WAN UNI to SD-WAN UNI. The Performance Monitoring metrics defined within this document are defined as ordered pairs between Monitoring Points. These Monitoring Points are located so that they provide a sub-set of a SD-WAN UNI to SD-WAN UNI view.

The requirements for Performance Monitoring of an SD-WAN Service and the related information that the SD-WAN Service Provider provides to the SD-WAN Subscriber are detailed within this standard.

In addition, Service Readiness Testing requirements for SD-WAN Services are defined within this standard. The Service Readiness Testing requirements are focused on the IP level because MEF 70.1 [13] defines an SD-WAN Service as an IP service where ‘the basic unit of transport at the SD-WAN UNI is an IP Packet’. Service Readiness Testing topics cover SD-WAN Service Readiness Measurement Point (SRMP), SD-WAN Test Methodologies, and the Service Readiness Testing report for an SD-WAN Service. Service Readiness Testing verifies that the SD-WAN Service is ready for the SD-WAN SP or Subscriber to implement the appropriate Policies and begin forwarding packets. It does not verify the operation of the Policies or SD-WAN Service Attributes. Instead, it is focused on verifying that the UCSs providing connectivity between the appropriate SD-WAN Edges are working correctly.

3 Terminology and Abbreviations

This section defines the terms used in this document. In many cases, the normative definitions to terms are found in other documents. In these cases, the third column is used to provide the reference that is controlling, in other MEF or external documents.

In addition, terms defined in MEF 48.1 [8], MEF 61.1 [10], MEF 66 [11], and MEF 70.1 [13] are included in this document by reference and are not repeated in the table below.

Term	Definition	Reference
AFMP	Application Flow Measurement Point	This document
AF Specification	Application Flow Specification	This document
Application Flow Measurement Point	A point where measurements of IP Packet performance are performed for Application Flows	This document
Application Monitoring Flow	A sequence of IP Packets that Ingress at an SD-WAN UNI or are directed towards an SD-WAN UNI from a UCS that match the same Application Flow Specification and have source IP Addresses in the same zoneName or are all in the zoneName <i>Internet</i> .	This document
CLEAR-TCA Window Threshold	The number of PM Metric Calculation Intervals, within the TCA Window Size, for which the PM Metric Value must be below the TCA Performance Threshold to generate a CLEAR-TCA, when using Stateful TCA Reporting.	This document
Collector Test Function	A logical function for counting and discarding received IP Packets, which can include test packets.	MEF 67 [12]
CTF	Collector Test Function	MEF 67 [12]
Damping Factor	A parameter with a numeric value that indicates the length of time to suppress new TCAs. Applicable to Stateless TCA Reporting only.	This document
Egress Measured Information Rate	The Measured Information Rate at the Egress of the SD-WAN UNI	This document
EMIR	Egress Measured Information Rate	This document
Generator Test Function	A logical function for generating and transmitting IP Packets which can include test packets.	MEF 67 [12]
GTF	Generator Test Function	MEF 67 [12]
IMIR	Ingress Measured Information Rate	This document
Information Rate	The rate, in bits per second, that IP Packets are received or transmitted to the Subscriber.	This document
Ingress Measured Information Rate	The Measured Information Rate at the Ingress of the SD-WAN UNI	This document

Term	Definition	Reference
IPDV Separation Time	The time between delay measurements that are used to calculate Inter-Packet Delay Variation	This document
Measured Information Rate	The Information Rate expressed in bits per second measured during a single PM Metric Calculation Interval for each Monitored Entity.	This document
Measurement IP Packet	Synthetic or Subscriber IP Packets that are used to measure performance.	This document
Monitored Entity	The SWVC UNI ordered pair or TVC End Point ordered pair that is being monitored.	This document
Monitored Entity Identifier	The identifier of a Monitored Entity. For TVCs it is the three-tuple (Ingress UCS End Point Identifier, Egress UCS End Point Identifier, UCS CoS Name). For Application Monitoring Flows it is the 4-tuple (AF Specification, zoneName, Ingress UNI Identifier, Egress UNI Identifier). For Application Flows, it is the 3-tuple (AF Specification, zoneName, UNI Identifier).	This document
Passive Monitoring	The monitoring of performance that does not use synthetic IP Packets or modify Subscriber IP Packets to perform measurements.	This document
Performance Monitoring	The collection of data concerning the performance of the Service. In this document, SD-WAN is the service for which Performance Monitoring is defined.	This document (derived from MEF 35.1 [6])
PM	Performance Monitoring	This document (derived from MEF 35.1 [6])
PM Implementation	An implementation of the requirements for SD-WAN PM in this standard.	This document
PM Metric	A metric that is measured or calculated as a part of Performance Monitoring.	This document
PM Metric Calculation Data Set	A set of PM Metric Values for a given PM Metric Calculation Interval.	This document
PM Metric Calculation Instance	The measurement of a given set of PM Metrics, using a given PM Metric Calculation Interval duration, over a given Monitored Entity.	This document
PM Metric Calculation Interval	The time interval over which one or more PM Metrics are calculated.	This document
PM Metric Calculation Profile	A profile that defines the PM Metrics that are calculated and the PM Metric Calculation Interval	This document
PM Metric Report	A report that contains the PM Metric Values for one or more PM Metric Calculation Intervals.	This document
PM Metric Report Interval	The time interval over which one or more PM Metric Values are reported to the Subscriber by the SD-WAN Service Provider.	This document

Term	Definition	Reference
SD-WAN Controller/Orchestrator	The entity that is responsible for managing/orchestrating the SD-WAN Service.	This document
SD-WAN PM	Performance Monitoring of the Application Monitoring Flows and TVCs in an SD-WAN Service.	This document
SD-WAN PM Implementation	An implementation that meets the requirements specified within this document for SD-WAN Performance Monitoring	This document
SD-WAN Test Function	An application for testing SD-WAN Service that resides on a SD-WAN Edge.	This document
Service Readiness Measurement Point	A reference point in the SD-WAN Service where events can be observed and measured during Service Readiness Testing. A Service Readiness Measurement Point contains both a Generator Test Function and a Collector Test Function.	This document derived from MEF 48.1 [8]
Service Readiness Testing	The testing that validates that the SD-WAN Service is ready to be configured with specific customer Policies. Testing verifies connectivity of ordered pairs of UCS End Points.	This document
Service Readiness Testing Parameters	The parameters for SRT tests that are agreed between the SD-WAN SP and the Subscriber.	This document
Service Readiness Testing Results	The results, which are included in the SRT report, of the SRT performed on each ordered pair of UCS End Points.	This document
SET-TCA Window Threshold	The number of PM Metric Calculation Intervals, within the TCA Window Size, for which the PM Metric Value must be at or above the TCA Performance Threshold to generate a SET TCA, when using Stateful TCA Reporting.	This document
SRMP	Service Readiness Measurement Point	This document
SRT	Service Readiness Testing	This document
SRT Parameters	Service Readiness Testing Parameters	This document
SRT Results	Service Readiness Testing Results	This document
Stateful TCA Reporting	A TCA reporting mechanism whereby a SET-TCA is generated when an alertable condition begins and a CLEAR-TCA is generated when it ends.	This document derived from MEF 35.1 [6]
Stateless TCA Reporting	A TCA reporting mechanism whereby TCAs are generated whenever an alertable condition is detected.	This document derived from MEF 35.1 [6]
TCA	Threshold Crossing Alert	This document derived from MEF 35.1 [6]
TCA Function	An instance of an implementation of Threshold Crossing Alerts for a given set of TCA Function parameters, a given PM Metric, and a given Monitored Entity.	This document

Term	Definition	Reference
TCA Performance Threshold	The PM Metric Value that is compared against, for each PM Metric Calculation Interval, when determining whether to generate a TCA.	This document
TCA Window Size	The sliding window of the number of consecutive PM Metric Calculation Intervals that are used to evaluate whether to generate a SET-TCA or CLEAR-TCA, when using Stateful TCA reporting.	This document
Threshold Crossing Alert	A notification message that is specific to a particular PM Metric, a particular set of TCA Function parameters and a particular Monitored Entity, and is generated when the PM Metric Value exceeds, equals, or falls below the TCA Performance Threshold.	This document derived from MEF 35.1 [6]
Tunnel Virtual Connection Measurement Point	The logical location between the SWVC End Point and the UCS UNI where measurements of TVC performance are performed.	This document
TVC MP	Tunnel Virtual Connection Measurement Point	This document
UTC	Coordinated Universal Time	IETF RFC 3339 [2]
zoneName	The name used to identify a specific Zone.	This document

Table 1 – Terminology and Abbreviations

It should be noted that throughout this document the term UNI is used to indicate an SD-WAN UNI unless otherwise noted. The term Service Provider (SP) is used to identify an SD-WAN Service Provider unless otherwise noted.

4 Compliance Levels

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in BCP 14 (RFC 2119 [1], RFC 8174 [3]) when, and only when, they appear in all capitals, as shown here. All key words must be in bold text.

Items that are **REQUIRED** (contain the words **MUST** or **MUST NOT**) are labeled as [**Rx**] for required. Items that are **RECOMMENDED** (contain the words **SHOULD** or **SHOULD NOT**) are labeled as [**Dx**] for desirable. Items that are **OPTIONAL** (contain the words **MAY** or **OPTIONAL**) are labeled as [**Ox**] for optional.

A paragraph preceded by [**CRa**]< specifies a conditional mandatory requirement that **MUST** be followed if the condition(s) following the "<" have been met. For example, "[**CR1**]<[D38]" indicates that Conditional Mandatory Requirement 1 must be followed if Desirable Requirement 38 has been met. A paragraph preceded by [**CDb**]< specifies a Conditional Desirable Requirement that **SHOULD** be followed if the condition(s) following the "<" have been met. A paragraph preceded by [**COc**]< specifies a Conditional Optional Requirement that **MAY** be followed if the condition(s) following the "<" have been met.

5 Conventions

5.1 Numerical Conventions

This document uses the prefix notation to indicate multiplier values as shown in Table 2.

Decimal		Binary	
Symbol	Value	Symbol	Value
k	10^3	Ki	2^{10}
M	10^6	Mi	2^{20}
G	10^9	Gi	2^{30}
T	10^{12}	Ti	2^{40}
P	10^{15}	Pi	2^{50}
E	10^{18}	Ei	2^{60}
Z	10^{21}	Zi	2^{70}
Y	10^{24}	Yi	2^{80}

Table 2 – Numerical Prefix Conventions

5.2 Diagram Conventions

The diagrams in this document have a number of components that appear frequently. These components are represented in a standard way as described in the following:

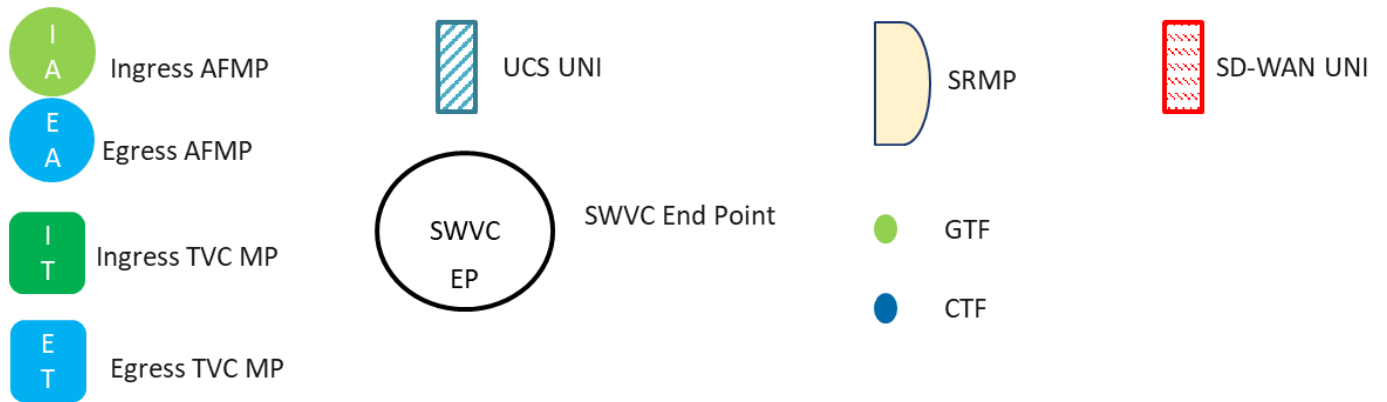


Figure 1 – Diagram Conventions

6 Introduction

For the SD-WAN Service standardized in MEF 70.1 [13], there is a need from both Service Providers and Subscribers to have a consistent method defined to monitor and report the performance of the SD-WAN Service. Additionally, there is a need to standardize the testing performed to verify that the SD-WAN Service is ready to be used by the Subscriber. SD-WAN Performance Monitoring (SD-WAN PM) provides a standardized method for monitoring the performance of the service. SD-WAN Service Readiness Testing (SD-WAN SRT) provides a standardized method for ensuring verify that the SD-WAN Service is ready to be used by the Subscriber.

This document defines an Application Monitoring Flow as a set of IP Packets that that match a given Ingress Application Flow and are directed towards a given Egress UNI. Note that IP Packets in a given Ingress Application Flow might be directed towards multiple Egress UNIs. An Application Flow may therefore contain multiple Application Monitoring Flows each directed towards a different Egress UNI.

Measurements of Application Monitoring Flows are performed between pairs of Application Flow Measurement Points which are located between SWVC End Points and UCS UNIs within an SD-WAN Edge as described in section 7.2.1.

This document describes the necessary requirements for the SD-WAN PM and SD-WAN SRT pertinent to an SD-WAN Service as described and defined in MEF 70.1 [13]. SD-WAN PM is described in Section 7 of this document. SD-WAN PM addresses the SD-WAN PM statistics and information the SD-WAN SP provides to the SD-WAN Subscriber, Performance Monitoring Metric Values per Application Monitoring Flow, Performance Monitoring Metric Values per pair of Tunnel Virtual Connection (TVC), ordered pair, Threshold Crossing Alerts (TCA) for the Performance Monitoring Metric Values.

Note: the term SD-WAN Performance Monitoring (PM) is limited to the functions and metrics defined within this document.

Note: SD-WAN PM includes monitoring TVCs but does not include monitoring UCS End Point pairs. This is because a TVC represent a forwarding relationship between two SD-WAN Edges. If a forwarding relationship does not exist between two SD-WAN Edges, then SD-WAN PM is not used between them.

PM Metric Values collected for SD-WAN Service are intended to be used by the SD-WAN SP and the Subscriber for many purposes including managing the service in real time. As an example of real time management, statistics on the number of packets transmitted and received on an SD-WAN UNI might be particularly important to a Subscriber who is trying to resolve an issue with communication to the location served by the UNI. If no packets are received from that location on the SD-WAN UNI at the SD-WAN Edge, the Subscriber may be able to quickly determine that the problem is with their equipment at that location. Similar metrics are collected for Application Monitoring Flows and TVCs. The SD-WAN Service Provider can use these to determine if

packets are passing over an Application Monitoring Flow and pass this information to the Subscriber e.g., via a Subscriber Portal.

Some PM Metrics can be used to make forwarding decisions within the SD-WAN Service. PM Metrics such as One-way Mean Packet Delay or One-way Packet Loss could be used to determine if a TVC is meeting the PERFORMANCE Policy criterion specified for an Application Flow. Forwarding of IP Packets of a given Application Flow over the TVC could be stopped if the TVC is not meeting the performance criteria.

The SD-WAN SRT is described in Section 8 of this document and addresses the following topics:

- SD-WAN Service Readiness Measurement Point (SRMP) functions and locations within the SD-WAN framework defined in MEF 70.1 [13]
- The test methodologies used for bringing an SD-WAN Service into readiness (i.e., ready for use by the Subscriber) including verifying connectivity of agreed to UCS End Point pairs included in an SWVC
- Definition of the SD-WAN SRT report including the attributes and metrics included in the report

Note: SRT is performed on UCS End Point pairs because there is no standard method for implementing TVCs. Since TVCs may be implemented before or after the SWVC is turned over to the Subscriber, UCS End Point pairs are tested as a part of SRT. In this way, connectivity of the UCS between UCS End Point pairs is verified so that if one or more TVCs is instantiated between the SD-WAN Edges, connectivity is known to have existed at the time of SRT.

Service Readiness Testing verifies that the SD-WAN Service is ready for the SD-WAN SP or Subscriber to implement the appropriate Policies and begin forwarding packets. It does not verify the operation of the Policies or the SD-WAN Service Attributes as defined in MEF 70.1 [13]. Instead, it is focused on verifying that the UCSs providing connectivity between the appropriate SD-WAN Edges are working correctly.

Areas that are not addressed within this document include the specific tools or implementations used to perform SD-WAN PM or SD-WAN SRT. In addition, there are a number of parameters that are agreed to by the SP and Subscriber that could be represented as Service Attributes but are not within the context of this document.

It should be noted that throughout this document the term UNI is used to indicate an SD-WAN UNI unless otherwise noted. The term Service Provider (SP) is used to identify an SD-WAN Service Provider unless otherwise noted.

7 SD-WAN Performance Monitoring

An SD-WAN SP may measure and report a variety of PM Metrics describing the performance of the various components of the SD-WAN Service. The tools used to measure SD-WAN performance are not defined in this document.

This document specifies the following aspects of SD-WAN PM:

- Performance Monitoring per Application Monitoring Flow defined in section 7.2.1
- Performance Monitoring per TVC in each direction defined in section 7.2.1
- PM Metrics for SD-WAN defined in section 7.2.2.2.
- Threshold Crossing Alerts (TCAs) defined in section 7.4

The Application Monitoring Flows that are monitored are agreed on by the SD-WAN Service Provider and the Subscriber. This may be a strict subset of all Application Monitoring Flows. Every TVC is monitored in both directions since each TVC represents a forwarding relationship between SD-WAN Edges, and it is assumed that SD-WAN PM is needed on all of these.

SD-WAN PM Metric Values are the results of PM measurements and calculations performed using IP PM Packets or other methods. A PM Implementation is an implementation of the requirements for SD-WAN PM in this standard.

Note: “Ingress” is used to describe packets received from the Subscriber at the SD-WAN UNI. “Egress” is used to describe packets that are transmitted towards the Subscriber at the SD-WAN UNI.

The SD-WAN PM Metrics specified in this document do not yield information about the performance of networks external to the SP-WAN Service Provider network, e.g., equipment used by the Subscriber to reach an SD-WAN UNI.

7.1 Performance Monitoring Framework

Performance Monitoring of an SD-WAN Service uses a different framework than more traditional Service OAM (SOAM) Performance Monitoring. SD-WAN PM starts with measurement intervals that are significantly shorter (10 seconds or less) than the traditional measurement intervals used in SOAM PM (5-15 minutes). PM measurements are made during the measurement interval and the metric value for that interval is calculated from the measurements. Interval metric values are used not only to measure the performance of Application Monitoring Flows or TVCs but can also be used to make IP Packet forwarding decisions when performance criteria are included in Policies.

The results of each SD-WAN measurement interval are reported based on the definition of the PM Metric Report detailed in section 7.3.2. The interval in which PM Metric Reports are generated and the PM Metrics included in the report are included in the PM Metric Report definition. PM Metric Reports can be used by the SD-WAN SP, and they may be shared with the Subscriber. If shared with the Subscriber, the Subscriber and SP agree to what is shared.

In addition to PM Metric Reports, one or more TCAs can be specified as part of SD-WAN Performance Monitoring. TCAs are useful for quickly identifying a service impacting degradation or fault. They may even be used to make IP Packet forwarding decisions.

In traditional SOAM PM the longer measurement intervals result in significantly more measurements taking place in the interval. Calculated values, as well as historical “bins””, are used to present the results of the measurement interval. Normally, PM statistics are not used to make forwarding decisions. Instead, they are used to measure compliance to Service Level Specifications (SLS).). It should be noted that there are no performance based SLSs defined for SD-WAN Service.

The PM Metrics, calculation, reporting, and the use of the PM metric values for Threshold Crossing Alerts are defined in the following sections of this document.

7.2 Performance Monitoring Metrics

This section describes the measurement methods and PM Metrics that are defined for SD-WAN Service. One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, and One-way Packet Loss Ratio are specified per agreed on Application Monitoring Flow and for each direction of each TVC. Measured Information Rate is specified for Ingress and Egress Application Flows located at different SD-WAN Edges.

One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, and One-way Packet Loss Ratio for IP Packets flowing from one SD-WAN UNI to another SD-WAN UNI over a Path are defined in MEF 70.1 [13] section 15. Corresponding definitions for PM Metrics, for IP Packets flowing over either an Application Monitoring Flow or TVC, are described in section 7.2.1.

7.2.1 Application Monitoring Flow and TVC PM Metrics

The definitions and requirements for measurements, calculation of PM Metrics, and the PM Metrics for Application Monitoring Flows and TVCs are contained in this section. The definitions used in support of these requirements are as follows:

- An Application Monitoring Flow is monitored at an Application Flow Measurement Point (AFMP).
- An Ingress AFMP is located after where any Policies are applied and the UCS UNI (see Figure 2).
- The Egress AFMP is located after the UCS UNI and before where any Policies are applied (see Figure 2).
- For One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, and One-way Packet Loss Ratio PM Metrics, measurements are performed between the ordered pair of Ingress AFMP adjacent to an SWVC End Point to an Egress AFMP adjacent to another SWVC End Point.

- For Measured Information Rate PM Metrics, measurements are performed before Ingress Policy is applied or after Egress Policy is applied (see Figure 3).

Note: Measured Information Rate is calculated based on byte counts collected by the Measured Information Rate byte counters. A Measured Information Rate of 0 Mb/s is considered a valid PM Metric Value. Implementations of Passive monitoring might result in PM Metrics that are reported as Undefined.

The location of Ingress and Egress AFMPs is shown in Figure 2.

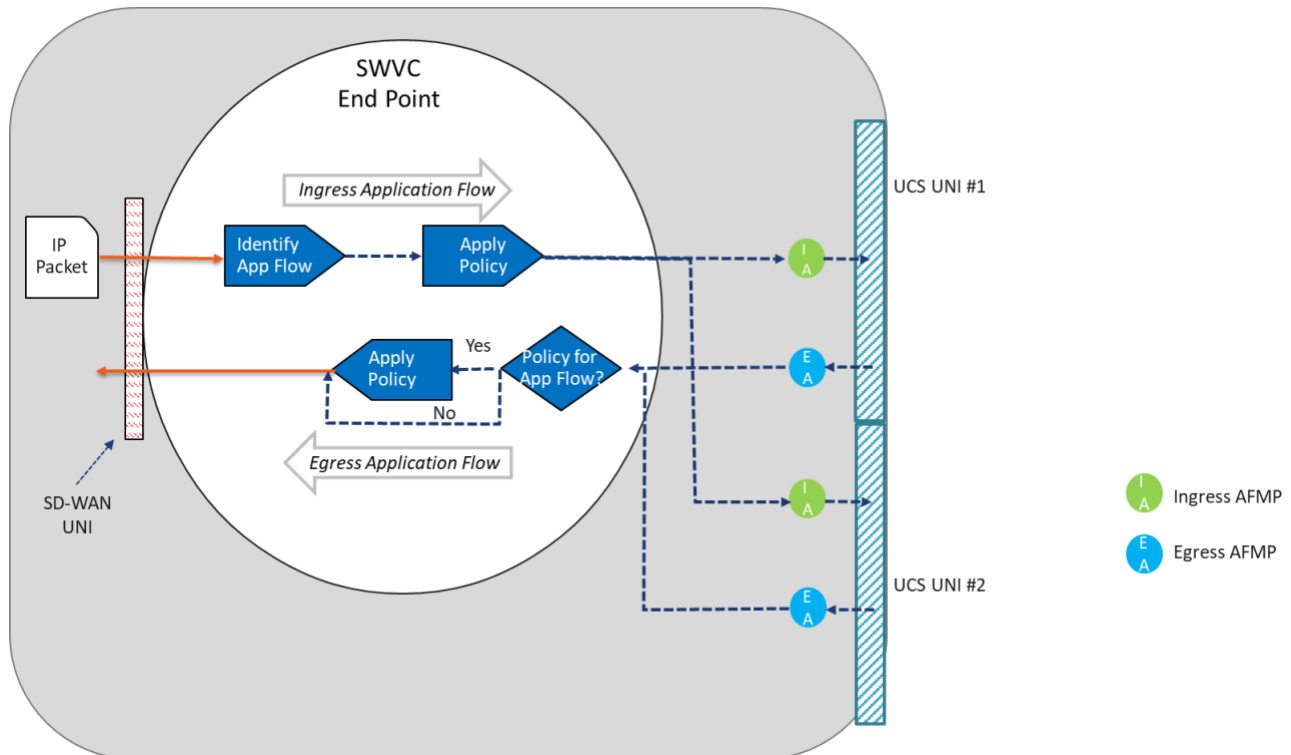


Figure 2 – AFMP Locations

The flow of Measurement IP Packets is from the Ingress AFMP at one SWVC End Point/UNI to the Egress AFMP at another SWVC End Point/UNI making an ordered pair. Measurement IP Packets are defined as synthetic or Subscriber IP Packets that are used to measure performance. This document does not specify the method or tool used to perform these measurements.

The location within the SWVC End Point of the Application Flow Ingress Measured Information Rate (IMIR) and Egress Measured Information Rate (EMIR) Byte Counters are shown in Figure 3.

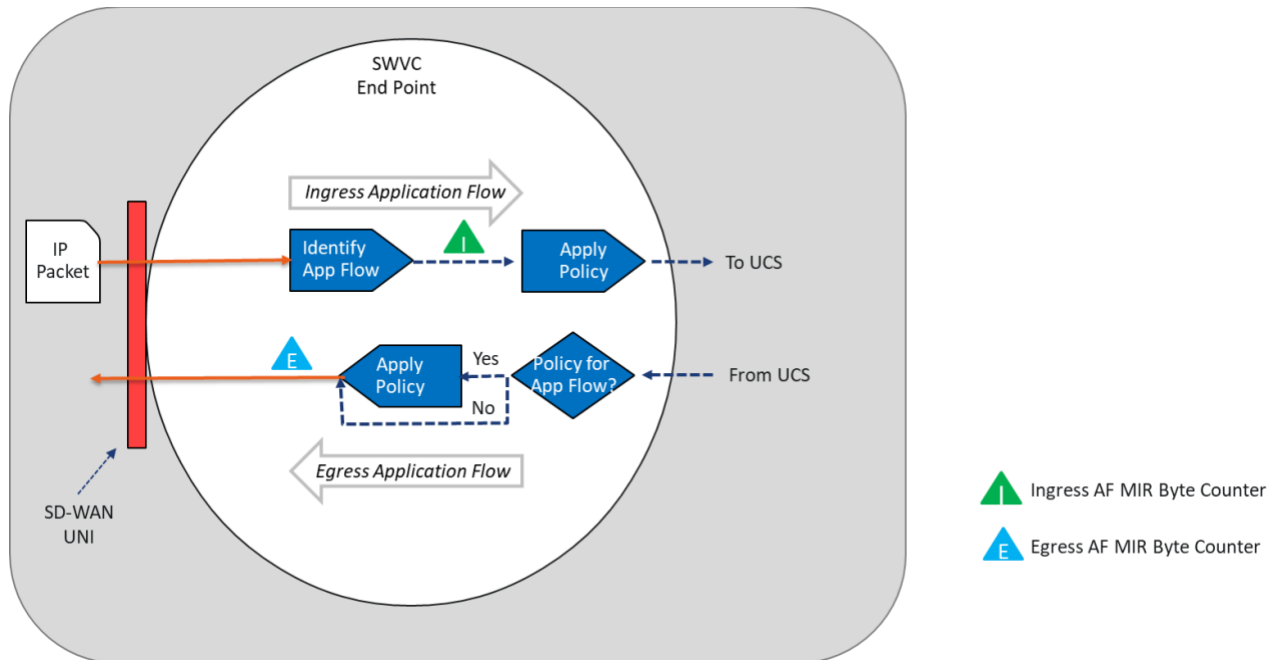


Figure 3 – Measured Information Rate Byte Counter Locations

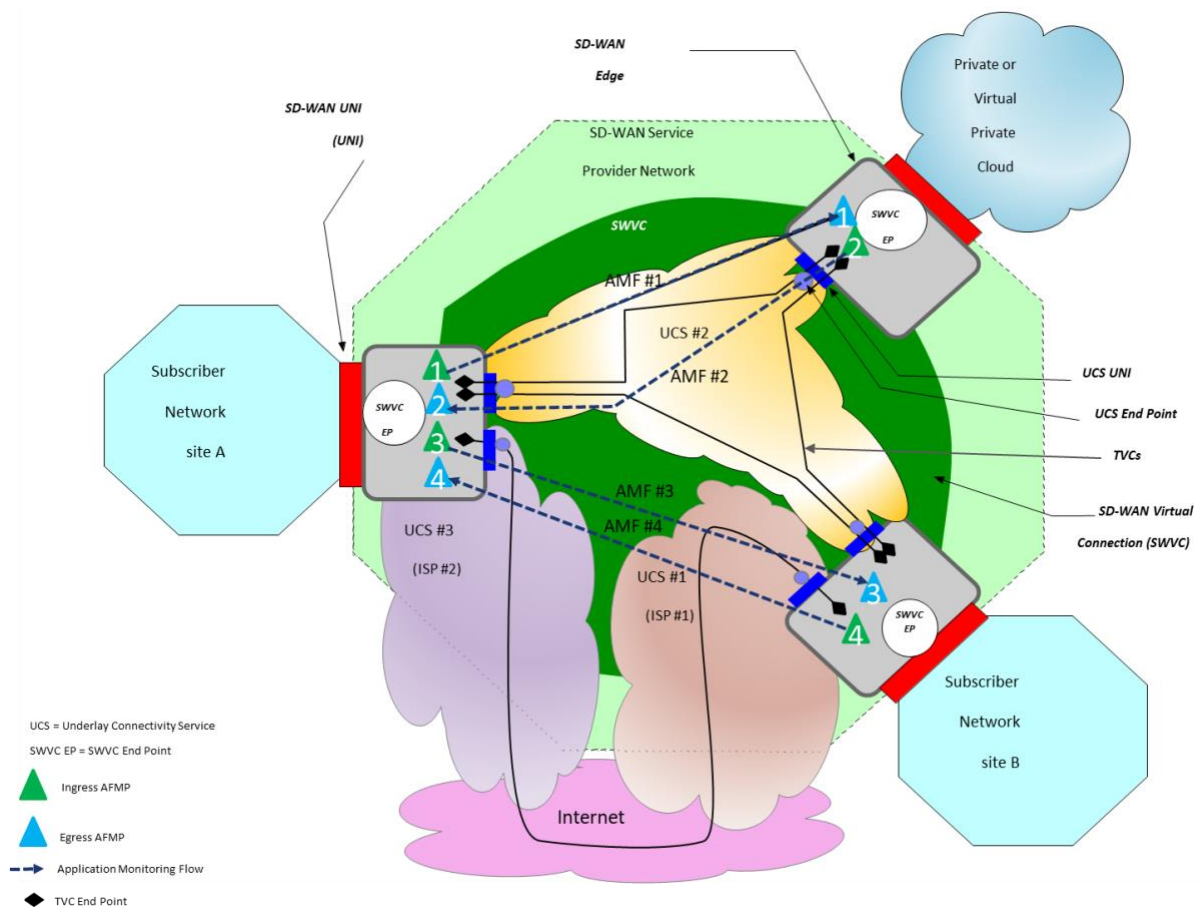


Figure 4 – Ingress AFMP to Egress AFMP ordered pair Example

Figure 4 shows four ordered pairs of AFMPs, identified as pair 1, pair 2, pair 3, and pair 4. The performance of IP Packets from the Ingress AFMP to the Egress AFMP is measured and PM Metric Values are calculated from the measurements. In the example, AFMP ordered pair 1 monitors Application Monitoring Flow 1, AFMP ordered pair 2 monitors Application Monitoring Flow 2, AFMP ordered pair 3 monitors Application Monitoring Flow 3, and AFMP ordered pair 4 monitor Application Monitoring Flow 4. For details on how one-way measurements are performed please see MEF 66 [11] section 8.3.

The following requirements apply to Application Monitoring Flows.

- [R1]** The Subscriber and SP **MUST** agree on the set of AF Specification, zoneName pairs to be monitored.
- [R2]** For each AF Specification, zoneName pair agreed on in [R1], the Subscriber and SP **MUST** agree to the set of PM Metrics to be monitored.
- [R3]** If metric values are being shared with the Subscriber, the Subscriber and SP **MUST** agree on the PM Metrics for which values are to be shared and the method by which they are shared.

The method used to share the metric values is outside the scope of this document.

An example of the agreed PM Metrics to be monitored is shown in Table 1.

AF Specification, zoneName	PM Metrics
Voice, Corporate	Packet Delay
	Inter-Packet Delay Variation
	Packet Loss Ratio
	Ingress Measured Information Rate
	Egress Measured Information Rate
Email, Corporate	Ingress Measured Information Rate
Video Streaming, Corporate	Packet Delay
	Ingress Measured Information Rate
	Egress Measured Information Rate

Table 1 – AF Specification, zoneName and PM Metric Agreement Example

- [R4]** The Subscriber and SP **MUST** agree on a single set of SD-WAN UNI ordered pairs for which PD, IPDV and PLR will be monitored for all AF Specification,

zoneName pairs for which those metrics have been agreed to be monitored per [R2].

An example is shown in Table 2.

SD-WAN UNI Identifier ordered pairs	
UNI 1	UNI 2
A	B
B	A
A	D
E	A
F	A
A	F

Table 2 – Monitored SD-WAN UNI ordered pair Agreement Example

Note that the combination of AF Spec and zoneName (from [R1]) and an ordered pair of UNIs (from [R4]) uniquely identifies an Application Monitoring Flow.

- [R5]** The Subscriber and SP **MUST** agree on a single set of SD-WAN UNIs for which Ingress Measured Information Rate and Egress Measured Information Rate will be monitored for all AF Specification, zoneName pairs for which those metrics have been agreed to be monitored per [R2].

An example is shown in Table 3.

SD-WAN UNI Identifiers
A
B
C
D
E
F
G

Table 3 – Monitored SD-WAN UNI Agreement Example

The SP performs measurements between Ingress and Egress AFMPs for each ordered pair of UNIs as shown in Table 2. For Measured Information Rate measurements, the SP uses Ingress and Egress Application Flow Measured Information Rate Byte Counters as shown in Figure 3.

Given the examples above, Table 4 shows an example of the PM Metrics that will be monitored for each AF Specification, zoneName and SD-WAN UNI ordered pair i.e., for each Application Monitoring Flow.

AF Specification, zoneName	SD-WAN UNI ordered pair	PM Metric
Voice, Corporate	A-B	One-way Mean Packet Delay One-way Inter- Packet Delay Variation One-way Packet Loss Ratio
Voice, Corporate	B-A	One-way Mean Packet Delay One-way Inter- Packet Delay Variation One-way Packet Loss Ratio
Voice, Corporate	A-D	One-way Mean Packet Delay One-way Inter- Packet Delay Variation One-way Packet Loss Ratio
Voice, Corporate	E-A	One-way Mean Packet Delay One-way Inter- Packet Delay Variation One-way Packet Loss Ratio
Voice, Corporate	A-F	One-way Mean Packet Delay One-way Inter- Packet Delay Variation One-way Packet Loss Ratio
Voice, Corporate	F-A	One-way Mean Packet Delay One-way Inter- Packet Delay Variation One-way Packet Loss Ratio
Video Streaming, Corporate	A-B	One-way Mean Packet Delay

AF Specification, zoneName	SD-WAN UNI ordered pair	PM Metric
Video Streaming, Corporate	B-A	One-way Mean Packet Delay
Video Streaming, Corporate	AAAA-D	One-way Mean Packet Delay
Video Streaming, Corporate	EEEE-A	One-way Mean Packet Delay
Video Streaming, Corporate	A-F	One-way Mean Packet Delay
Video Streaming, Corporate	F-A	One-way Mean Packet Delay

Table 4 – AF Specification, zoneName, ordered pair, and PM Metric Example

Given the examples above, Table 5 shows an example of the PM Metrics that will be monitored for each AF Specification, zoneName pair and a given SD-WAN UNI i.e., for each Application Flow.

AF Specification, zoneName	SD-WAN UNI	PM Metric
Voice, Corporate	A	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	B	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	C	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	D	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	E	Ingress Measured Information Rate Egress Measured Information Rate

AF Specification, zoneName	SD-WAN UNI	PM Metric
Voice, Corporate	F	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	G	Ingress Measured Information Rate Egress Measured Information Rate
Email, Corporate	A	Ingress Measured Information Rate
Email, Corporate	B	Ingress Measured Information Rate
Email, Corporate	C	Ingress Measured Information Rate
Email, Corporate	D	Ingress Measured Information Rate
Email, Corporate	E	Ingress Measured Information Rate
Email, Corporate	F	Ingress Measured Information Rate
Email, Corporate	G	Ingress Measured Information Rate
Video Streaming, Corporate	A	Ingress Measured Information Rate Egress Measured Information Rate

AF Specification, zoneName	SD-WAN UNI	PM Metric
Video Streaming, Corporate	B	Ingress Measured Information Rate Egress Measured Information Rate
Video Streaming, Corporate	C	Ingress Measured Information Rate Egress Measured Information Rate
Video Streaming, Corporate	D	Ingress Measured Information Rate Egress Measured Information Rate
Video Streaming, Corporate	E	Ingress Measured Information Rate Egress Measured Information Rate
Video Streaming, Corporate	F	Ingress Measured Information Rate Egress Measured Information Rate
Video Streaming, Corporate	G	Ingress Measured Information Rate Egress Measured Information Rate

Table 5 – AF Specification zoneName, UNI Identifier, and PM Metric Overview Example

- [D1]** For each AF Specification, zoneName $\langle afs, z \rangle$ pair that is agreed to be monitored, One-way Mean Packet Delay and One-way Packet Loss Ratio **SHOULD** be included in the list of agreed PM Metrics.
- [R6]** For a given pair of Application Flow Specification and zoneName, $\langle afs, z \rangle$, agreed per [R1], and a given ordered pair of UNIs $\langle u1, u2 \rangle$ agreed per [R4], the SP **MUST** monitor and report performance for the Application Monitoring Flow $\langle iaf, u2 \rangle$, where iaf is the Ingress Application Monitoring Flow identified by the 3-tuple $\langle afs, z, u1 \rangle$, for each of the Performance Monitoring Metrics that were agreed for the pair $\langle afs, z \rangle$ per [R2]:
- One-way Mean Packet Delay

- One-way Mean Inter-Packet Delay Variation
- One-way Packet Loss Ratio

[R7] For a given pair of Application Flow Specification and zoneName, $\langle afs, z \rangle$, agreed per [R1], and a given UNI ul agreed per [R5], the SP **MUST** monitor and report performance for the Ingress and Egress Application Flows identified by the 3-tuple $\langle afs, z, ul \rangle$, for each of the Performance Monitoring Metrics that were agreed for the pair $\langle afs, z \rangle$ per [R2]:

- Application Flow Ingress Measured Information Rate (for the Ingress Application Flow)
- Application Flow Egress Measured Information Rate (for the Egress Application Flow)

The definitions and requirements for measurements, calculation of PM Metrics, and the PM Metrics for TVCs are detailed below. The definitions used in support of these requirements are as follows:

- A TVC is defined as a forwarding relationship between two SD-WAN Edges. PM is performed unidirectionally on a TVC.
- A TVC is monitored at a TVC Measurement Point (TVC MP).
- An Ingress TVC MP is located after where the ingress Policy is enforced and before the corresponding UCS UNI.
- An Egress TVC MP is located after the corresponding UCS UNI and before where the egress Policy is enforced.
- For One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, and One-way Packet Loss Ratio PM Metrics, measurements are performed between the ordered pair of Ingress TVC MP located on one SD-WAN Edge and Egress TVC MP located on another SD-WAN Edge as shown in Figure 5.

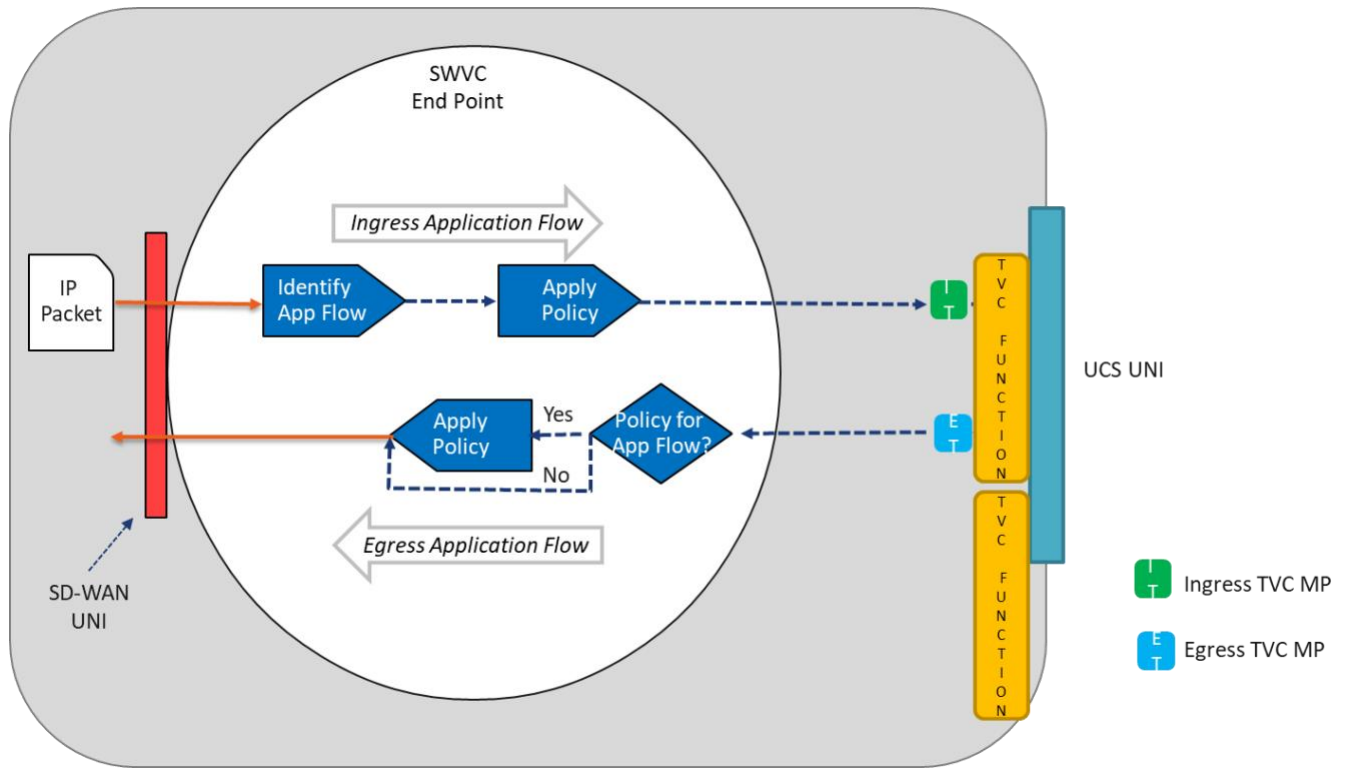


Figure 5 – TVC MP Locations

The flow of Measurement IP Packets is from the Ingress TVC MP located at the Ingress UCS UNI to the Egress TVC MP located at Egress UCS UNI.

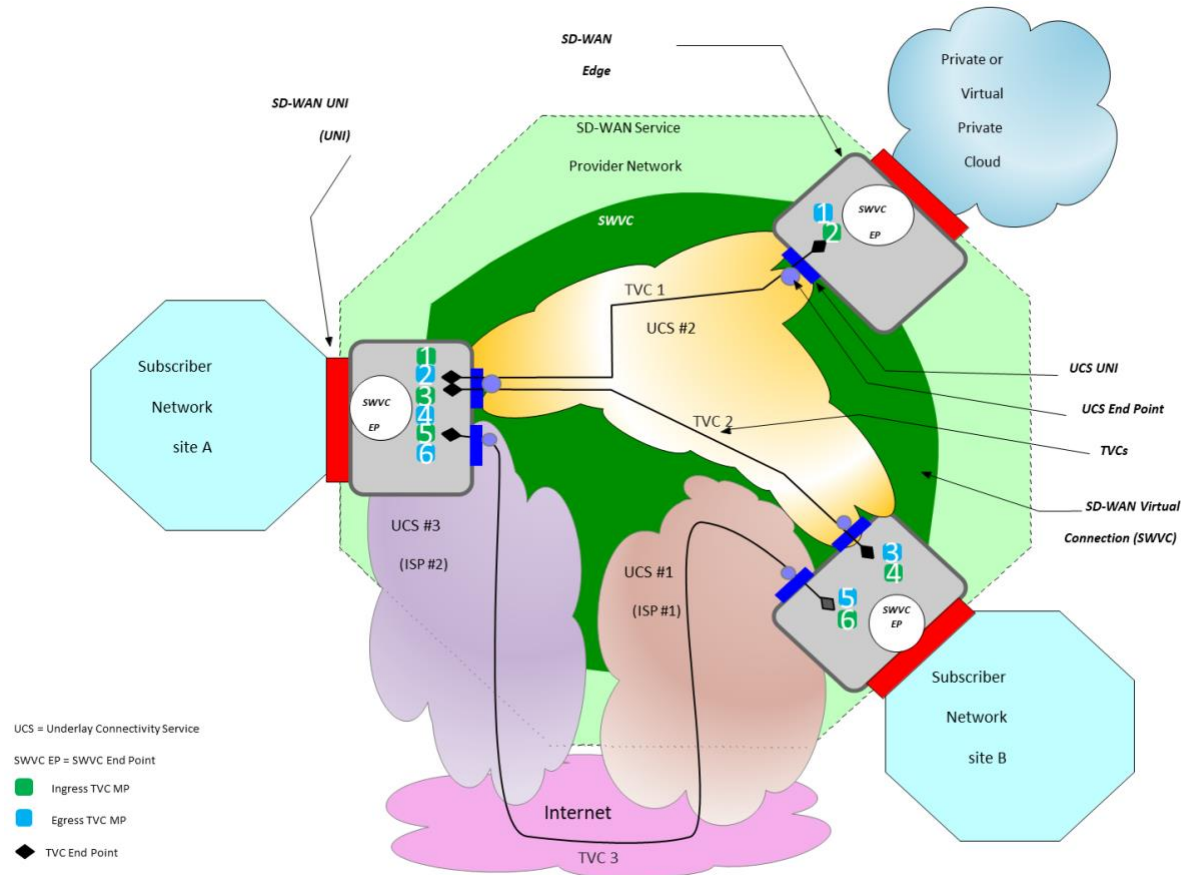


Figure 6 – Ingress TVC MP to Egress TVC MP ordered pair Example

Figure 6 shows six TVC MP ordered pairs, identified as 1, 2, 3, 4, 5, and 6, that are monitoring three TVCs. Each ordered pair is made up of an Ingress TVC MP and an Egress TVC MP with the same number. IP Packets are exchanged between the Ingress and Egress TVC MPs, PM Metric measurements are performed using these packets, and PM Metric Values are calculated from the measurements. Two ordered pairs of TVC MPs are used to monitor the performance of each TVC. In the example, TVC MP ordered pairs 1 and 2 monitor TVC 1, TVC MP ordered pairs 3 and 4 monitor TVC 2, and TVC MP ordered pairs 5 and 6 monitor TVC 3.

[R8] The SD-WAN SP **MUST** monitor all TVCs in an SWVC in both directions for the following PM Metrics:

- One-way Mean Packet Delay
- One-way Mean Inter-Packet Delay Variation
- One-way Packet Loss Ratio

All TVCs within a given SWVC are monitored using the same duration as specified in [R16]. This allows forwarding decisions based on performance. As an example, if the performance of two TVCs is worse on one TVC than the other, the better performing TVC can be selected. If different

durations were used, a forwarding decision could not be made based on performance over the same period of time.

Note: TVCs are monitored to inform the Subscriber of the performance of the overlay. In addition, Application Flow forwarding decisions based on performance criteria can be made by the Service Provider.

7.2.2 PM Configuration

A PM Metric configuration is defined in this section. This includes the Monitored Entity, the PM Metric List, the PM Metric Calculation Interval, and the PM Metric Calculation Instance. The normative text and requirements that are intended to be met by any SD-WAN PM Implementation.

7.2.2.1 Monitored Entity

A Monitored Entity is one of the following:

1. An Application Monitoring Flow flowing between a pair of UNIs, identified by the 4-tuple *<AF Specification, zoneName, Ingress UNI Identifier, Egress UNI Identifier>* used for at least one of One-way Mean Packet Delay, One-way Inter-Packet Delay Variation and One-way Packet Loss Ratio of an Application Monitoring Flow
2. An Application Flow in a given direction at a UNI, identified by the 4-tuple *<AF Specification, zoneName, UNI Identifier>* - used for Ingress MIR and Egress MIR of an Application Flow
3. A TVC in a given direction, identified by the 3 tuple *<Ingress UCS End Point Identifier, Egress UCS End Point Identifier, UCS CoS Name>* - used for One-way Mean Packet Delay, One-way Inter-Packet Delay Variation and One-way Packet Loss Ratio of the TVC.

Note: Not all UCS support CoS Names. In the event that a UCS does not support CoS (such as an L1 UCS) Null is used to indicate the UCS CoS Name rather than a non-empty list.

7.2.2.2 PM Metric List

For each Monitored Entity, the Subscriber and SP agree to a list of PM metrics that will be monitored. This list contains one or more PM Metrics as specified within this document.

[R9] The PM Metric List used for Application Monitoring Flows **MUST** contain all or a non-empty subset of the following PM Metrics:

- One-way Mean Packet Delay
- One-way Inter-Packet Delay Variation
- One-way Packet Loss Ratio

[R10] A PM Metric List used for Application Flows **MUST** contain one or both of the following PM Metrics:

- Ingress Measured Information Rate
- Egress Measured Information Rate

[R11] The PM Metric List for TVCs **MUST** contain the following PM Metrics:

- One-way Mean Packet Delay
- One-way Inter-Packet Delay Variation
- One-way Packet Loss Ratio

7.2.2.2.1 PM Metrics for Application Monitoring Flows

The definitions for the PM Metrics for Application Monitoring Flows are detailed below.

The One-way Packet Delay for an IP Packet belonging to a given Ingress Application Flow that flows from an Ingress AFMP at a given Ingress UNI for the Ingress Application Flow to an Egress AFMP at a given Egress UNI is defined as:

- The time elapsed from the transmission of the first bit of the IP Packet at the Ingress AFMP until the reception of the last bit of the first corresponding IP Packet at the Egress AFMP. If the IP Packet is erroneously duplicated as it traverses the network, the delay is based on the first copy that is delivered.

The One-way Mean Packet Delay for an Application Monitoring Flow and for a PM Metric Calculation Interval is defined as:

- Let $\Delta = \{\delta_1, \delta_2, \delta_3, \dots, \delta_n\}$ represent the set of the One-way Packet Delays of the n IP Packets associated with the Application Monitoring Flow from Ingress AFMP to Egress AFMP transmitted during the PM Metric Calculation Interval. Then the One-way Mean Packet Delay for the Application Monitoring Flow over that interval is the arithmetic mean of the values $\delta_1 \dots \delta_n$.
- If Δ is *null* (no measurement values) then the One-way Packet Delay for the PM Metric Calculation Interval is reported as *Undefined*.

One-way Mean Inter-Packet Delay Variation for an Application Monitoring Flow and for a PM Metric Calculation Interval is defined as:

- Let $\Delta = \{\kappa_1, \kappa_2, \kappa_3, \dots, \kappa_n\}$ represent the set of the One-way Packet Delays of the n IP Packets associated with the Application Monitoring Flow that were transmitted from the Ingress AFMP during the PM Metric Calculation Interval to the Egress AFMP during a time interval whose duration is the value of the PM Metric Calculation Interval duration.

- Let Δ' = the set of all pairs of elements $\{\kappa_m, \kappa_n\}$ in Δ such that $n > m$ and the difference in the time of transmission at the Ingress AFMP of packets n and m equals the *arrival interval* in the value of SWVC Performance Time Intervals Service Attribute as defined in MEF 70.1 [13].
- If Δ' is *null* (no measurement values) then the One-way Mean Inter-Packet Delay Variation for the PM Metric Calculation Interval is reported as *Undefined*. Otherwise, let v_{mn} be the absolute value of the difference in One-way Packet Delay for each pair, $\{\kappa_m, \kappa_n\}$ in Δ' , i.e., $v_{mn} = |\kappa_m - \kappa_n|$. Then the One-way Mean Inter-Packet Delay Variation for the Application Flow over the PM Metric Calculation Interval is the arithmetic mean of the values v_{mn} for each element in Δ' .

Packet Loss Ratio for Application Flow and a PM Metric Calculation Interval is defined as:

- Let s represent the total number of IP Packets in the Application Monitoring Flow a that are transmitted from the Ingress AFMP towards the Egress AFMP during the PM Metric Calculation Interval.
- Let r represent the total number of IP Packets received from the Ingress AFMP by the Egress AFMP for the Application Monitoring Flow that were sent during the same interval. Then the One-way Packet Loss Ratio over that interval for the Application Monitoring Flow is defined as follows:
 - If $s=0$, then the One-way Packet Loss Ratio is 0.¹
 - If $s>0$, then the One-way Packet Loss Ratio is $(s-r)/s$

[R12] The One-way Packet Loss Ratio **MUST** be represented as a percentage.

Bits included in the Measured Information Rate measurements are those that belong to an IP Packet.

7.2.2.2.2 PM Metrics for Application Flows

Application Flow Ingress Measured Information Rate for a given Ingress Application Flow and a given PM Metric Calculation Interval is defined as:

- The total number of bits in IP Packets that pass a given Ingress MIR measurement point during that interval, divided by the duration of the interval (which in this case is the PM Metric Calculation Interval duration for Ingress MIR for Application Flows).

Application Flow Egress Measured Information Rate for a given Egress Application Flow and a given PM Metric Calculation Interval is defined as:

¹ In theory, this can only happen if Performance is measured on user-data. If Performance is measured on synthetic traffic, then there should be at least period times the synthetic-rate qualified packets.

- The total number of bits in IP Packets that pass a given Egress MIR measurement point during that interval, divided by the duration of the interval (which in this case is the PM Metric Calculation Interval duration for Egress MIR for Application Flows).

Note: IP Packet is defined in MEF 61.1 [10] as “Either an IPv4 Packet or an IPv6 Packet, from the start of the IP Version field to the end of the IP data field”. The AF MIR byte counters are located so they only count IP Packets.

Measured Information Rate, in bits per second, is used throughout this document and is calculated as shown below:

$$\text{Measured Information Rate} = \frac{\text{count of bytes during PMCI duration} \times 8}{\text{PMCI duration}}$$

- [R13]** An SD-WAN PM Implementation **MUST** count the bytes of all ingress IP Packets associated with an Application Flow before any Ingress Policy, as specified in MEF 70.1 [13] and shown in Figure 3, is applied when calculating Application Flow Ingress Measured Information Rate.
- [R14]** An SD-WAN PM Implementation **MUST** count the bytes of all egress IP Packets associated with an Application Flow after any Egress Policy, as specified in MEF 70.1 [13] and shown in Figure 3, is applied when calculating Application Flow Egress Measured Information Rate.

7.2.2.2.3 PM Metrics for TVCs

The definitions for the PM Metrics for TVCs are detailed below. The impact of Values on forwarding decisions is beyond the scope of this document.

The One-way Packet Delay for an IP Packet belonging to a given TVC that flows from an Ingress TVC MP to an egress TVC MP is defined as:

- The time elapsed from the transmission of the first bit of the IP Packet at the Ingress TVC MP until the reception of the last bit of the first corresponding IP Packet at the Egress TVC MP. If the IP Packet is erroneously duplicated as it traverses the network, the delay is based on the first copy that is delivered.

One-way Mean Packet Delay for packets flowing in one direction over a given TVC from an ingress TVC MP to an egress TVC MP, and for a given PM Metric Calculation Interval is defined as:

- Let $\Delta = \{\delta_1, \delta_2, \delta_3, \dots, \delta_n\}$ represent the One-way Packet Delay of the n IP Packets sent over the TVC from Ingress TVC MP to Egress TVC MP that were transmitted from the ingress TVC MP during the PM Metric Calculation Interval during a time interval whose duration is the value of the PM Metric Calculation Interval duration. Then the One-way Mean Packet Delay for TVC t over that interval is the arithmetic mean of the values $\delta_1 \dots \delta_n$.

- If Δ is *null* (no measurement values) then the One-way Mean Packet Delay for the PM Metric Calculation Interval is *Undefined*.

One-way Mean Inter-Packet Delay Variation for packets flowing in one direction over a given TVC from an ingress TVC MP to an egress TVC MP, and for a given PM Metric Calculation Interval is defined as:

- Let $\Delta = \{\delta_1, \delta_2, \delta_3, \dots, \delta_n\}$ represent the set of the One-way Packet Delays of the n IP Packets sent over the TVC from Ingress TVC MP to Egress TVC MP that were transmitted from the ingress TVC MP during the PM Metric Calculation Interval.
- Let $\Delta' =$ the set of all pairs of elements $\{\delta_o, \delta_p\}$ in Δ such that $p > o$ and the difference in the arrival time of the first bit at the Ingress TVC MP of packets p and o equals the duration of the *arrivalInterval*.
- If Δ' is *null*, then the One-way Mean Packet Delay Variation for the PM Metric Calculation Interval is *Undefined*. Otherwise, let v_{op} be the absolute value of the difference in One-way Packet Delay for each pair, $\{\delta_o, \delta_p\}$ in Δ' , i.e., $v_{op} = |\delta_o - \delta_p|$. Then the One-way Mean Packet Delay Variation for t over that interval is the arithmetic mean of the values v_{op} for each element in Δ' .

One-way Packet Loss Ratio for packets flowing in one direction over a given a TVC from an ingress TVC MP to an egress TVC MP, and for a given PM Metric Calculation Interval duration is defined as:

- Let s represent the total number of IP Packets sent over the TVC from Ingress TVC MP to Egress TVC MP during the PM Metric Calculation Interval.
- Let r represent the total number of IP Packets received from Ingress TVC MP by the Egress TVC MP on the TVC that were sent during the same interval. Then the One-way Packet Loss Ratio over that PM Metric Calculation Interval for the TVC is defined as follows:
 - If $s=0$, then the One-way Packet Loss Ratio is 0.²
 - If $s>0$, then the One-way Packet Loss Ratio is $(s-r)/s$

The One-way Packet Loss Ratio is represented as a percentage.

7.2.2.3 PM Metric Calculation Duration

The PM Metric Calculation Duration is the period of time over which PM measurements are taken and PM Metrics are calculated from the measurements. The Subscriber and SP agree to the value of the PM Metric Calculation Duration. A PM Metric Calculation Duration is used for Monitored Entities. All TVCs use the same PM Metric Calculation Interval duration.

² In theory, this can only happen if Performance is measured on user-data. If Performance is measured on synthetic traffic, then there should be at least period*synthetic-rate qualified packets.

- [R15]** A PM Metric **MUST** use the same PM Metric Calculation Interval duration for all Application Monitoring Flow/SWVC End Point ordered pairs within the same SWVC.

[R15] means that there can be a maximum of three PM Metric Calculation Intervals (one per PM Metric) for Application Monitoring Flows and a maximum of two PM Metric Calculation Intervals (one per PM Metric) for Application Flows.

- [O1]** Different PM Metric Calculation Duration values **MAY** be used for different Application Flow Specification, zoneName pairs.

- [R16]** All three PM Metrics specified in [R8] that are calculated for TVCs **MUST** use the same PM Metric Calculation Interval duration for all TVC MP ordered pairs within the same SWVC.

[R16] means that there can be a maximum of one PM Metric Calculation Interval for TVCs that are monitored.

- [R17]** The SD-WAN SP and the Subscriber **MUST** agree on the value of each PM Metric Calculation Interval duration.

- [R18]** The PM Metric Calculation Interval duration **MUST** be less than or equal to 10000 milliseconds.

The PM Metric Calculation Intervals repeat based on the PM Metric Calculation Interval duration defined.

7.2.2.4 PM Metric Calculation Instance

The attributes of a PM Metric Calculation Instance are defined within this section. A PM Metric Calculation Instance is the measurement of a given set of PM Metrics, using a given PM Metric Calculation Interval duration, over a given Monitored Entity. The PM Metric Calculation Instance references the following:

- Monitored Entity
- PM Metric List
- PM Metric Calculation Interval duration

An example of a PM Metric Calculation Instance would include a Monitored Entity of *Ingress UCS EP Identifier 5West, Egress UCS EP Identifier 5East, UCS CoS Name Gold*, a PM Metric List of *One-way Mean Packet Delay, One-way Inter-Packet Delay Variation, and One-way Packet Loss Ratio*, and a PM Metric Calculation Interval of *100 milliseconds*. *UCS CoS Name* setting depends on the underlying UCS technology. When the UCS is Layer 3, the CoS Name is as described in MEF 61.1 [10]. When the UCS is Layer 2, the CoS Name is as described in MEF 10.4 [6]. When the UCS is Layer 1, the CoS Name is simply a placeholder. Measurements

included within the calculation are performed at some rate that is less than or equal to the PM Metric Calculation Interval.

Attribute	Value	Value Constraints	Comments
PM Metric Calculation Instance Identifier	An identifier that is unique within the SP for the PM Metric Calculation Instance	1-53 characters	The format of the identifier is not defined in this standard and is left to the implementation.
SWVC Identifier	The identifier of the SWVC	String	See MEF 70.1 [13] section 7.16 for value constraints.
Monitored Entity Identifier	Identifier for the Monitored Entity	1-53 characters	The format of the identifier is not defined in this standard and is left to the implementation.
PM Metric List	The list of PM Metrics that are include in the PM Metric Calculation Instance	For Application Monitoring Flows One or more of: <ul style="list-style-type: none"> • One-way Mean Packet Delay • One-way Inter-Packet Delay Variation • One-way Packet Loss Ratio For Application Flows one or more of the following: <ul style="list-style-type: none"> • Ingress Measured Information Rate • Egress Measured Information Rate For TVCs the following: <ul style="list-style-type: none"> • One-way Mean Packet Delay • One-way Inter-Packet Delay Variation 	

Attribute	Value	Value Constraints	Comments
		<ul style="list-style-type: none"> One-way Packet Loss Ratio 	
PM Metric Calculation Interval	The duration in milliseconds of the PM Metric Calculation Interval	A real number from 1-10000	

Table 6 – PM Metric Calculation Instance Attributes

[R19] An SD-WAN PM Implementation **MUST** use the attributes specified in Table 6 for the PM Metric Calculation Instance.

The value of the attributes in Table 6 are agreed to by the Subscriber and SP.

An SD-WAN PM Implementation may use additional attributes for the PM Metric Calculation Instance. These attributes are outside the scope of this document.

[R20] A given Monitored Entity **MUST NOT** appear in more than one PM Metric Calculation Instance for a given PM Metric.

The method used to delete an PM Metric Calculation Instance and stop monitoring is beyond the scope of this document. Whether the PM Metric Values can be retrieved after a PM Metric Calculation Instance is deleted is left to the implementation.

7.3 PM Metric Calculation and Reporting

This section describes the configuration, calculation, and reporting of PM Metrics by the SD-WAN SP. Figure 7 illustrates the high-level sub-processes that are included in PM Metric Calculation and Reporting.

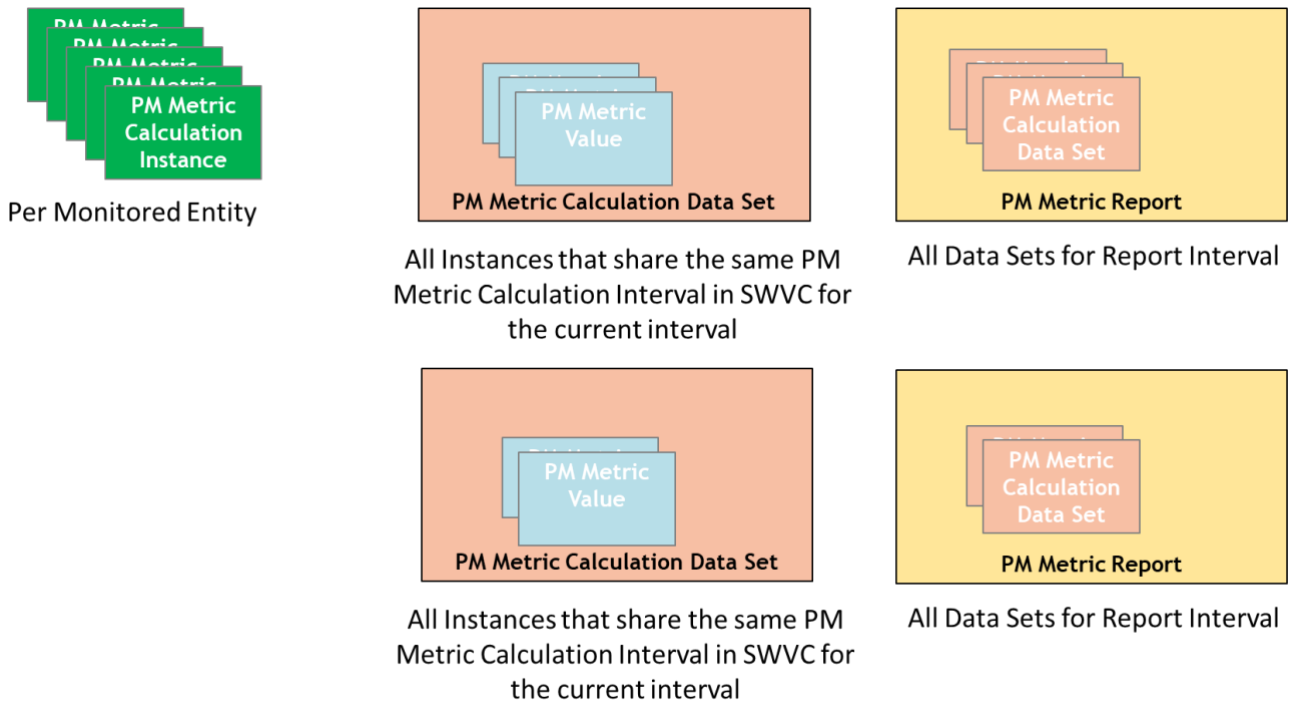


Figure 7 – PM Metric Calculation and Reporting Example

Once a PM Metric Calculation Instance is created and PM measurements for the Monitored Entity have begun, a PM Metric Value is calculated for each of the PM Metric Calculation Intervals. How the value is calculated is specified in section 7.2.

To simplify reporting of PM Metric Values, the PM Metric Calculation Intervals are aligned among different Monitored Entities i.e., they use the same duration. The initial PM Metric Value for a given Monitored Entity may be calculated over a time period less than the specified duration to allow the Monitored Entity's PM Metric Calculation Intervals to align.

The PM Metric Calculation Data Set is defined as the set of PM Metric Values for a given PM Metric Calculation Interval. This PM Metric Value is included in a PM Metric Calculation Data Set that contains the PM Metric Values for a given PM Metric Calculation Interval. There is a sequence of PM Metric Calculation Data Sets for each PM Metric Calculation Interval duration. In the example above, there are two different PM Metric Calculation Interval durations being used which results in two PM Metric Calculation Data Sets being created. A PM Metric Calculation Data Set contains the following attributes:

- PM Metric Calculation Data Set Start Time
- PM Metric Calculation Data Set End Time
- List of PM Metric Calculation Instance Identifiers
- List of pairs per Instance consisting of PM Metric and PM Metric Value

These attributes are defined in section 7.3.1.

One or more PM Metric Calculation Data Sets are reported in a PM Metric Report. The PM Metric Report is defined with a PM Metric Report Interval that specifies the start and end times of the PM Metric Report. PM Metric Reports are generated at the completion of each PM Metric Report Interval. This process repeats for each PM Metric Report Interval. The PM Metric Report contains the following attributes:

- PM Metric Report Interval Start Time
- PM Metric Report Interval End Time
- List of PM Metric Calculation Data Sets whose end time is within the PM Metric Report Interval.

The PM Metric Report duration included in the PM Metric Report is defined as the duration of the interval, in milliseconds.

The above attributes are defined in section 7.3.2.

The relationship between PM Metric Calculation Data Sets and PM Metric Reports are shown below.

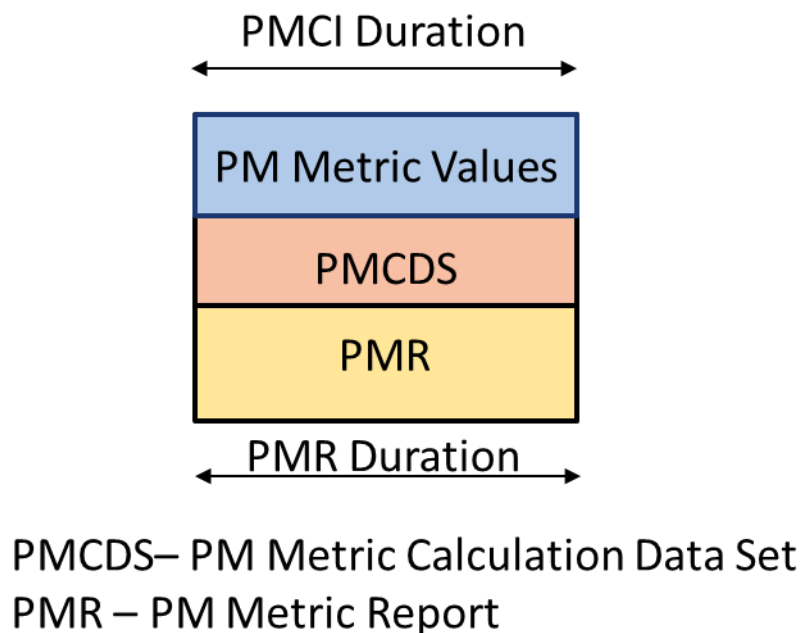


Figure 8 – PM Metric Calculation Data Set Duration Equals the PM Metric Report Duration Example

Figure 8 shows the PM Metric Calculation Interval duration and the PM Metric Report Interval duration are equal and with the start times aligned. All PM Metric Calculation Instances that use that duration are included in the report.

The PM Metric Calculation Interval duration and the PM Metric Report Duration do not have to be equal. Figure 9 shows an example of this.

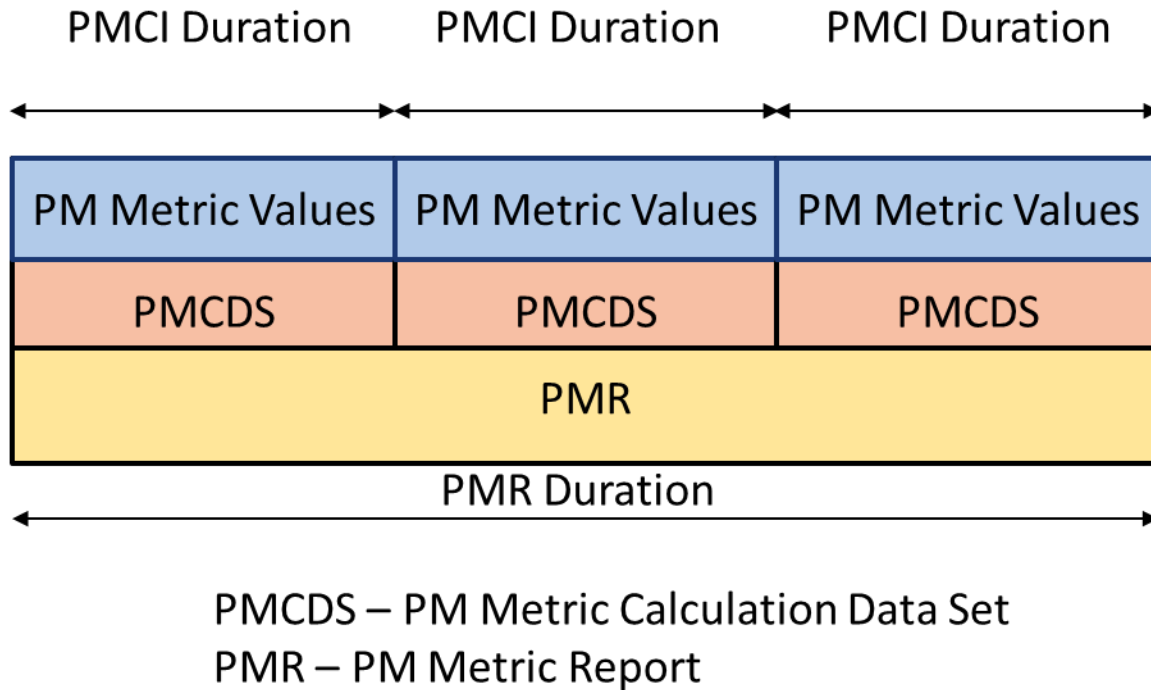


Figure 9 – PM Metric Report Interval Greater Than PM Metric Calculation Data Set Duration

In this example, there are three PM Metric Calculation Data Sets in each PM Metric Report.

- [R21]** The PM Metric Report Duration **MUST** be an integer multiple of each of the PM Metric Calculation Interval durations for the PM Metric Data Sets that are included in the PM Metric Report.

7.3.1 PM Metric Calculation Data Set

The attributes of a PM Metric Calculation Data Set are defined within this section.

Attribute	Value	Value Constraint	Comments
PM Metric Calculation Data Set Start Time	The Date and Time that the PM Metric Calculation Data Set started measurement and calculation for the PM Metric Calculation Instances in the data set	UTC	
PM Metric Calculation Data Set Duration	The length of time of the PM Metric Calculation Data Set Duration	≤ 10000 ms	
List of PM Metric Calculation Instance Identifiers	a list of identifiers of the PM Metric Calculation Instances	1-53 characters for each item in the list	
List of PM Metrics and PM Metric Values	A list of PM Metrics and PM Metric Values for each PM Metric Calculation Instance Identifier contained within the data set.		This list contains values for PM Metric as specified in Table 6.

Table 7 – PM Metric Calculation Data Set Attributes

[R22] An SD-WAN PM Implementation **MUST** include the attributes defined in Table 7 when reporting a data set.

[R23] For each Monitored Entity and each PM Metric included in the PM Metric Calculation Data Set, the PM Metric Calculation Data Set **MUST** include exactly one PM Metric Value.

7.3.2 PM Metric Report

The attributes of a PM Metric Report are defined in Table 8.

Attribute	Value	Value Constraint	Comments
PM Metric Report Interval Start Time	Time that the oldest PM Metric Calculation Data Set	Point in time in UTC	

Attribute	Value	Value Constraint	Comments
	contained in the PM Metric Report starts.		
PM Metric Report Interval duration	The duration, in ms, of the PM Metric Report.	Length of Time	
List of PM Metric Calculation Data Sets	One or more PM Metric Calculation Data Sets including all PM Metrics and PM Metric Values per PM Metric Calculation Instance		

Table 8 – PM Metric Report Attributes

- [R24]** The SD-WAN SP **MUST** provide a PM Metric Report that contains the attributes specified in Table 8.
- [R25]** If an aggregated value of a PM Metric over the PM Metric Report Interval is reported, the aggregated value of the PM Metric **MUST** be provided in addition to the per PM Metric Calculation Data Set values.
- [R26]** A SD-WAN PM implementation **MUST** generate a PM Metric Report at the completion of each PM Metric Report Interval.

Note: there may be some difference in time between the PM Metric Report Interval Duration and the time that the report is generated. This is due to the time it takes the SD-WAN PM implementation to collect the results for the last PM Metric Calculation Data Set included in the PM Metric Report Interval.

- [D2]** The SD-WAN SP **SHOULD** support the ability to store the PM Metric Values for a period of time as agreed to by the SD-WAN SP and Subscriber and for the PM Metric Values to be available for the Subscriber to retrieve.

The period of time agreed to by the SD-WAN SP and the Subscriber needs to conform to any legal requirements. Those requirements and the duration of the period of time are outside the scope of this document.

There are several mechanisms that can be used by the SD-WAN SP to enable the Subscriber to view these PM Metric Values. These mechanisms can range from a dashboard that refreshes PM Metric Values at some interval greater than or equal to the PM Metric Calculation Interval to informing the Subscriber of a performance degradation using Threshold Crossing Alerts (TCAs) as defined in section 7.4. This document does not mandate how the PM Metric Values are reported to the Subscriber.

7.4 Threshold Crossing Alerts

Threshold Crossing Alerts (TCAs) can be configured for certain PM Metrics (One-way Mean Packet Delay, One-way Inter-Packet delay Variation, One-way Packet Loss Ratio) and used to detect when performance is degraded beyond a given pre-configured level. From an SD-WAN perspective, TCAs can be used by the SD-WAN SP, and/or the SD-WAN SP can convey TCAs to the Subscriber.

Within this document the term TCA Function describes the implementation of Threshold Crossing Alerts. A TCA Function is the application of a given set of TCA parameters for a given PM Metric to a given PM Metric Calculation Instance.

- [R27] If TCAs are shared with the Subscriber, the threshold values and other parameters of the TCAs **MUST** be agreed on by the SP and Subscriber.
- [R28] If TCAs are shared with the Subscriber, the method by which the Subscriber is informed of the TCA **MUST** be agreed on by the SP and Subscriber.

Note that a TCA does not have to be established for every PM Metric that is being calculated. Reporting of TCAs is described in section 0.

There are two types of TCA reporting, *Stateful* and *Stateless*. Stateful TCA reporting is used to possibly reduce the total number of TCAs that are generated. The intent of Stateful TCAs is to provide a notification when a degradation is first encountered, followed by another when the degradation is resolved. A Stateful TCA Function uses the following TCA Function Parameters defined in section 7.4.1 to determine if a TCA should be set or cleared:

- TCA Performance Threshold Value
- SET TCA Window Threshold
- CLEAR TCA Window Threshold
- TCA Window Size

This contrasts with Stateless TCA reporting, in which TCAs are generated when a degradation is first encountered, for each PM Metric Calculation Interval that meets or exceeds the TCA Performance Threshold Value for as long as the degradation lasts subject to the Damping Factor. A TCA Function that uses Stateless TCA reporting uses the following TCA parameters defined in section 0 to determine if a TCA should be set:

- TCA Performance Threshold Value
- Damping Factor (desirable not mandatory)

An issue that can exist when using a Stateless TCA function is that a degradation that exists for more than one PM Metric Calculation Interval results in multiple TCAs being declared. A degradation that exists for several PM Metric Calculation Intervals can result in a flood of TCAs being generated possibly overwhelming SD-WAN SP alarm management systems. To avoid this, an optional attribute for Stateless TCA functions is defined; this is known as the Damping Factor. The Damping Factor is a method used to suppress new TCAs. The Damping Factor defines a number of consecutive PM Metric Calculation Intervals where the PM Metric Value is equal to or greater

than the TCA Performance Threshold Value and the new TCAs are suppressed for that number of PM Metric Calculation Intervals.

The TCA Function parameters that are agreed on by the Subscriber and the SP for each TCA Function are shown below:

- PM Metric
- Monitored Entity
- TCA Reporting Type (*Stateful* or *Stateless*)
- Damping Factor *enabled* or *disabled* if the TCA Reporting Type is Stateless
- Damping Factor Value (number of PM Metric Calculation Intervals for which TCAs are dampened)
- The TCA Performance Threshold Value is configured for a particular PM Metric. A TCA can be generated when the PM Metric Value for a set of PM Metric Calculation Intervals violates, the configured TCA Performance Threshold Value.
- When the TCA Reporting Type is Stateful, the TCA Window Thresholds define the number of PM Metric Calculation Intervals where the PM Metric Value is either does not violate (CLEAR), or violates (SET), the TCA Performance Threshold Value.
- The TCA Window Size defines the sliding window of the number of consecutive PM Metric Calculation Intervals over which the TCA is assessed.

Within this document a single TCA Function is specific to a given PM Metric Calculation Instance and a given PM Metric, but it might result in many TCAs. The TCA Function parameters defined for a given TCA Function are defined for a given Monitored Entity (i.e., Ingress Application Flow, Egress UNI) ordered pair or a (TVC, direction) pair. A Monitored Entity can have multiple TCA Functions associated with it, for the same or different PM Metrics.

Note: the use of TCAs to make forwarding decisions is beyond the scope of this document.

Stateful and Stateless TCA Reporting are explained in the following sections.

7.4.1 Stateful Threshold Crossing Alert Reporting

When using Stateful TCA reporting, each TCA Function has two configured TCA Window Threshold Values: a SET-TCA threshold and a CLEAR-TCA threshold.

The TCA Function also has an internal state, which may be ‘set’ or ‘clear’. The algorithm is informally described in the following two paragraphs and the flowchart. The algorithm is defined in the following requirements in section

The TCA Function begins in the 'clear' state and stays in the 'clear' state for the first PM Metric Calculation Interval. A SET-TCA is generated when the PM Metric Value is equal to or greater than the TCA Performance Threshold Value for the number of PM Metric Calculation Intervals defined by the SET-TCA Window Threshold out of the number of PM Metric Calculation Intervals defined by TCA Window Size. The TCA Function is then considered to be in a 'set' state, and no further SET-TCAs are generated in this state if the condition that triggered the TCA continues.

The TCA Function moves from the 'set' state to the 'clear' state when the PM Metric Value is less than the TCA Performance Threshold Value for the number of PM Metric Calculation Intervals defined by CLEAR-TCA Window Threshold out of the number of PM Metric Calculation Intervals defined by TCA Window Size. A CLEAR-TCA is generated, and the Stateful TCA Function returns to the 'clear' state. Thus, each SET-TCA can only be followed by a single CLEAR-TCA.

The Stateful TCA algorithm is shown in the flowchart in Figure 10. The flowchart does not specify a specific implementation of the algorithm. Any implementation can be used that exhibits the behavior regarding the reporting of TCAs

For intervals 1,2, ...

$St(n)$ = TCA Function value for interval n (Set or Clear)

W = Set TCA Window Size, S = SET-TCA Window Threshold,

C = CLEAR-TCA Window Threshold

$P(j)$ = PM Metric Value for interval j

$M(n)$ = number of intervals with $P(j) \geq$ TCA Performance Threshold for $j = n, n-1, n-2, \dots, n-W+1$

$N(n)$ = number of intervals with $P(j) <$ TCA Performance Threshold for $j = n, n-1, n-2, \dots, n-W+1$

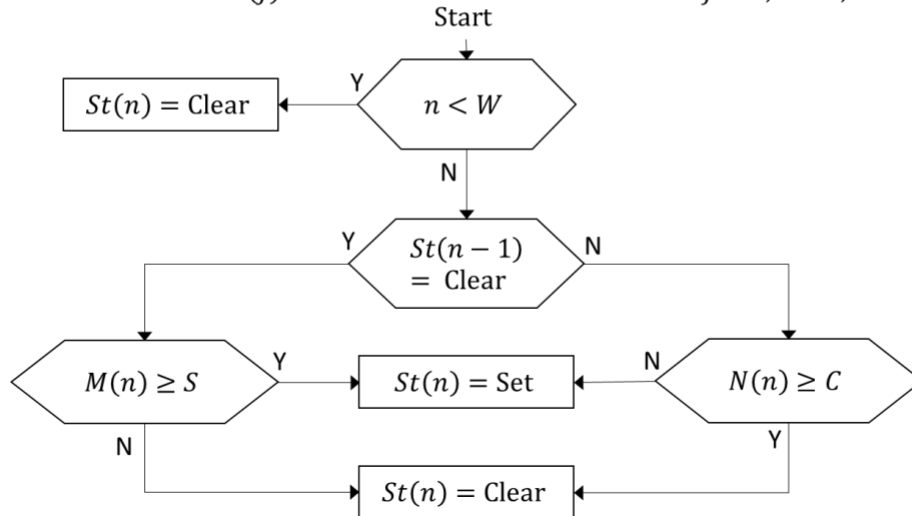


Figure 10 – Stateful TCA Flowchart

The algorithm is formally defined in the following requirements in section 7.4.3.3:

- [CR5]<
- [CR6]<
- [CR7]<

- [CR8]<
- [CR9]<
- [CR10]<
- [CR11]<
- [CR12]<
- [CR13]<
- [CR14]<
- [CR15]<

Figure 11 shows an example of Stateful TCA reporting. The sum of the SET-TCA Window Threshold and the CLEAR-TCA Window Threshold must be greater than the TCA Window Size (as required per [CR12]< in section 0). This is mandated within this document to avoid a condition where the SET-TCA and CLEAR-TCA criteria are met at the same time.

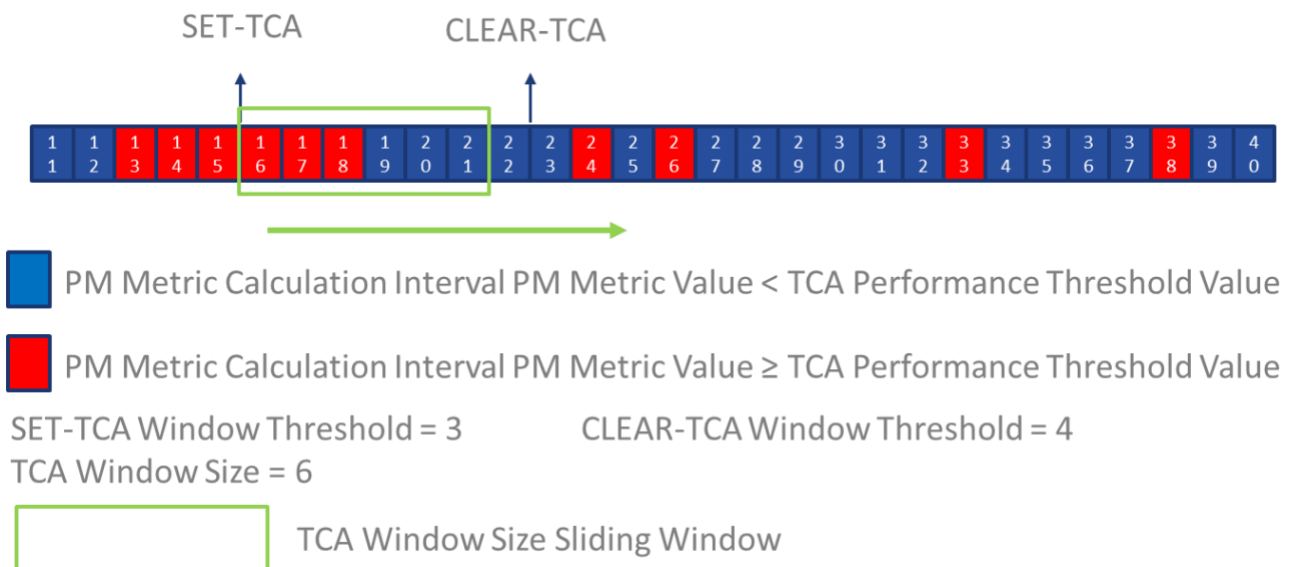


Figure 11 – Stateful TCA

Figure 11 shows an example of a Stateful TCA Function SET-TCA and CLEAR-TCA. In this example, the TCA Window Size is set to 6 and moves to the right in the figure with the completion of each PM Metric Calculation Interval. The SET-TCA Window Threshold is set to 3 and the CLEAR-TCA Window Threshold is set to 4.

A Stateful SET-TCA is generated when the PM Metric Value is equal to or greater than TCA Performance Threshold Value for SET in PM Metric Calculation Intervals 13, 14, and 15 (i.e., 3

out of 5 which means that the SET-TCA Window Threshold out of TCA Window Size criterion has been met).

The Stateful CLEAR-TCA is generated when the PM Metric Value is less than TCA Performance Threshold Value for CLEAR in PM Metric Calculation Intervals 19, 20, 21, and 22 (i.e., 4 out of 6 CLEAR-TCA Window Threshold out of TCA Window Size criterion met).

While the PM Metric Value is equal to or greater than the TCA Performance Threshold Value in PM Metric Calculation Intervals 24, 26, 33, and 38, the SET-TCA Window Threshold and TCA Window Size criterion is not met so another SET-TCA is not generated.

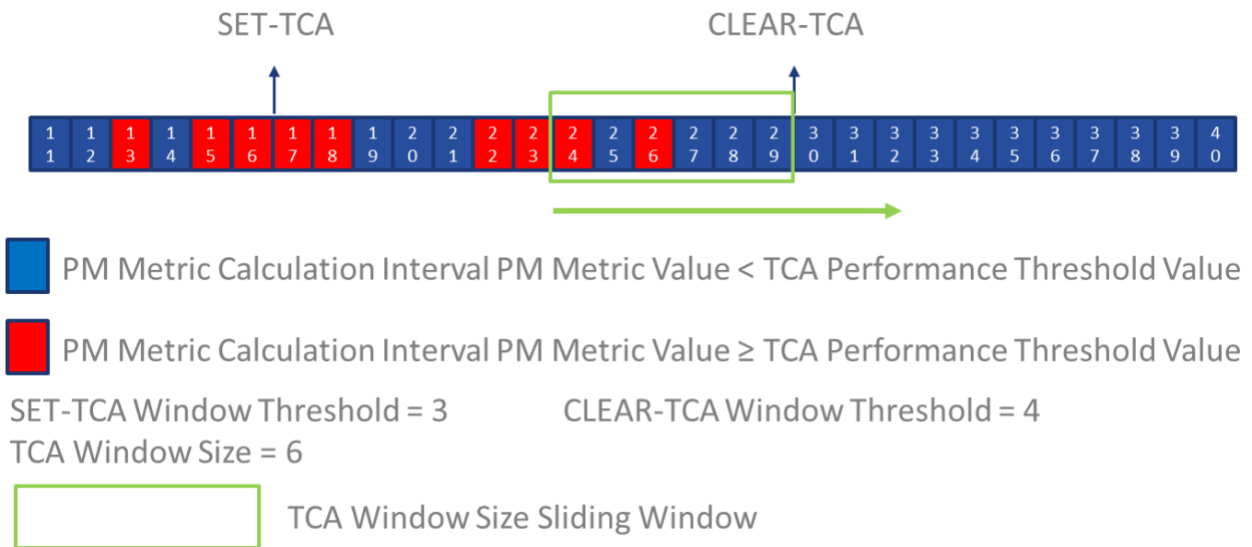


Figure 12 – Stateful TCA without Consecutive PM Metric Calculation Intervals ≥ TCA Performance Threshold Value Example

Figure 12 shows an example of a SET-TCA and CLEAR-TCA being generated even if consecutive PM Metric Calculation Intervals do not meet the criteria for SET-TCA or CLEAR-TCA. The PM Metric Value is equal to or greater than the TCA Performance Threshold Value in PM Metric Calculation Intervals 13, 15, and 16 meeting the SET-TCA Window Threshold and TCA Window Size criteria and a SET TCA is generated at the end of PM Metric Calculation Interval 6. A CLEAR-TCA is generated after the completion of PM Metric Calculation Interval 29 because the CLEAR-TCA Window Threshold and TCA Window Size criteria have been met, even though the PM Metric Value in PM Metric Calculation Interval 16 exceeded the TCA Performance Threshold Value.

7.4.2 Stateless Threshold Crossing Alert Reporting

The Stateless TCA Function treats each PM Metric Calculation Interval separately. The TCA Window Thresholds and TCA Window Size are not used.

The algorithm is informally described in the following paragraph and the flowchart. The algorithm is defined in the following requirements in section

When using Stateless TCA reporting, each TCA Function has a single configured TCA Performance Threshold Value. When a PM Metric Value in a PM Metric Calculation Interval is equal to or greater than the TCA Performance Threshold Value for a PM Metric Calculation Interval, a TCA is generated. There is no corresponding CLEAR.

The algorithm for Stateless TCAs without the Damping Factor is informally shown in Figure 13.

For intervals $n = 1, 2, \dots$

$$V(n) = \begin{cases} 1 & \text{if there is a threshold violation during interval } n \\ 0 & \text{otherwise} \end{cases}$$

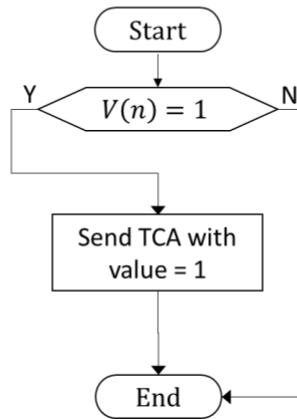


Figure 13 – Stateless TCA w/o Damping Factor Flowchart

The algorithm is formally defined in the following requirements in section 7.4.3.3:

- [CR16]<
- [CR17]<

The algorithm for Stateless TCA with Damping Factor is informally described in Figure 14. The flowchart does not specify a specific implementation of the algorithm. Any implementation can be used that exhibits the behavior regarding the reporting of TCAs

The algorithm is formally defined in the following requirements in section 7.4.3.3:

- [CD2]<
- [CR18]<
- [CR19]<
- [CR20]<

- [CR21]<
- [CR22]<
- [CR23]<
- [CD3]<
- [CR25]<
- [CR26]<
- [CR27]<

For intervals $n = 1, 2, \dots$

$d =$ Damping Factor ≥ 2 ,

c is a counter

$A = \begin{cases} \text{True} & \text{if Damping Algorithm is Active} \\ \text{False} & \text{otherwise} \end{cases} \quad (A = \text{False for } n = 1)$

$S(n) = \sum_{j=n-d+1}^n V(j), \quad V(n) = \begin{cases} 1 & \text{if there is a threshold violation during interval } n \\ 0 & \text{otherwise} \end{cases}$

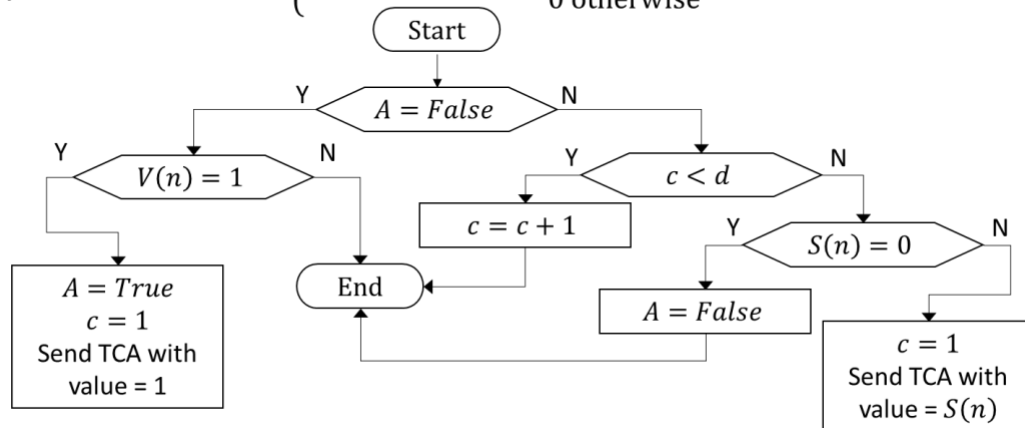


Figure 14 – Stateless TCA w/Damping Factor Flowchart

Figure 15 shows an example of a Stateless TCA without the Damping Factor.

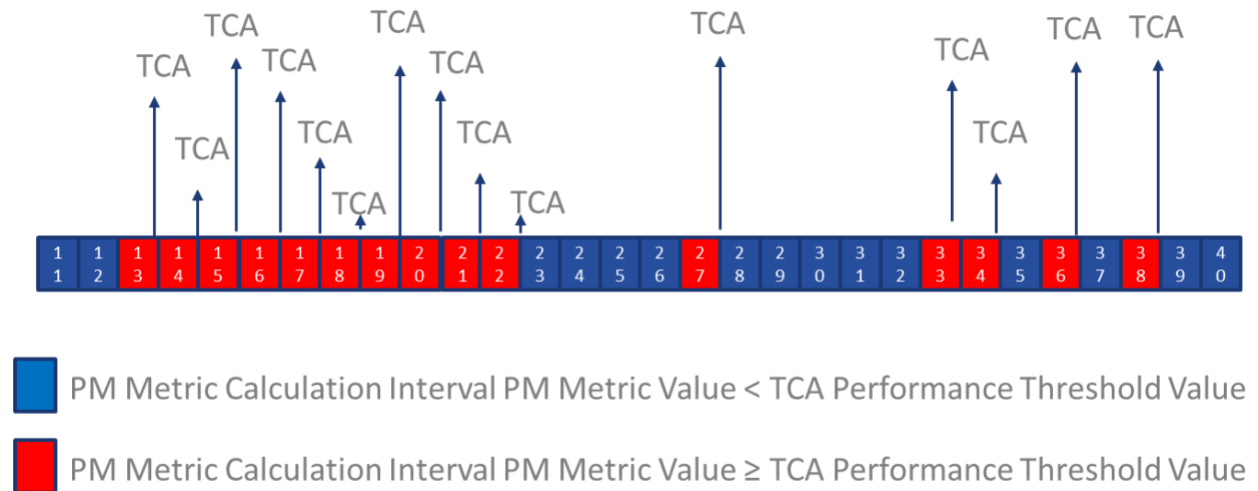


Figure 15 – Stateless TCA Example

As shown in Figure 15, multiple TCAs are generated when a degraded condition exists for more than one PM Metric Calculation Interval, one per PM Metric Calculation Interval that meets the TCA Function Parameters, when Stateless TCAs Reporting is used. In this example, the TCA criterion is met in PM Metric Calculation Interval 13-22, 27, 33, 34, 36, and 38. TCAs are generated in each of these PM Metric Calculation Intervals.

To avoid generating a TCA per PM Metric Calculation Interval when the TCA Function Parameters is met for multiple PM Metric Calculation Intervals, the Damping Factor is used. The impact of the Damping Factor is shown in Figure 16.

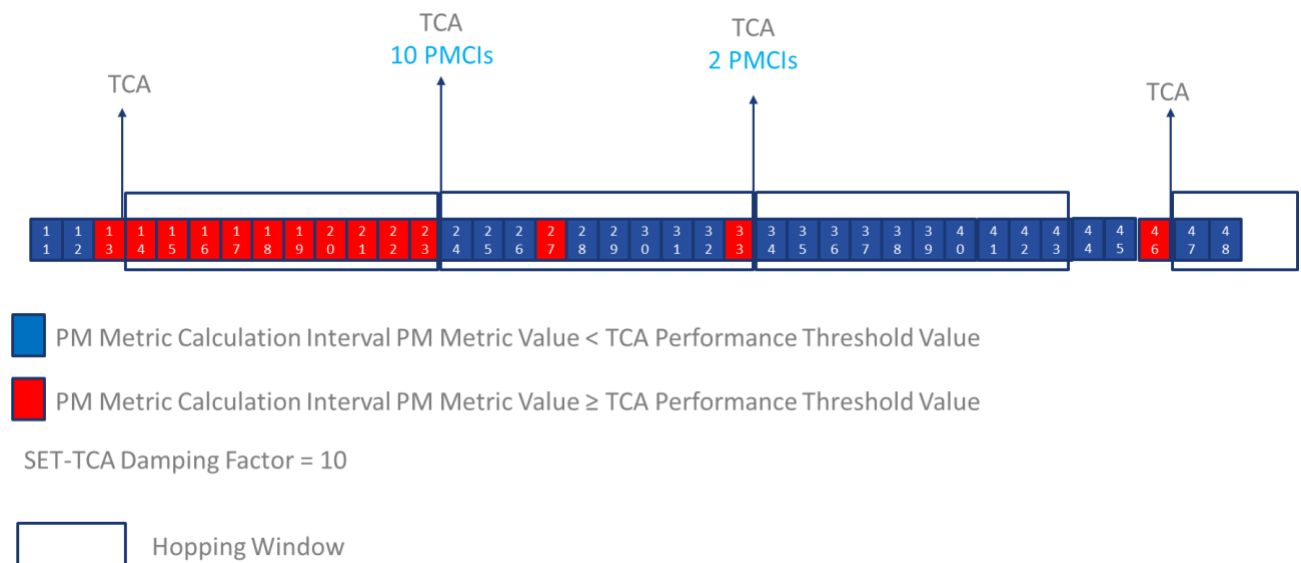


Figure 16 – Stateless TCA with Damping Factor Example

As shown in Figure 16, rather than multiple TCAs being generated, the TCAs are suppressed for the number of consecutive PM Metric Calculation Intervals defined by the SET-TCA Damping Factor. In the example, the PM Metric Value is equal to or greater than the TCA Performance Threshold Value in PM Metric Calculation Interval 13. An initial TCA is generated at the end of the PM Metric Calculation Interval, and the Damping algorithm is activated. A hopping window that is equal to the number of PM Metric Calculation Intervals specified by the Damping Factor starts at the completion of PM Metric Calculation Interval 13. Moving forward 10 PM Metric Calculation Intervals (Damping Factor value) another TCA is generated at the completion of PM Metric Calculation Interval 23 since at least one PM Metric Calculation Interval within the hopping window had a PM Metric Value equal to or greater than the TCA Performance Threshold Value. This TCA includes the count of PM Metric Calculation Intervals in the hopping window that had a PM Metric Value equal to or greater than the TCA Performance Threshold Value (shown as 10 PMCIs in blue text in Figure 16). Moving forward another 10 PM Metric Calculation Intervals to PM Metric Calculation Interval 33, a TCA is generated since at least one PM Metric Calculation Interval within the hopping window had a PM Metric Value equal to or greater than the TCA Performance Threshold Value. This TCA includes the count of PM Metric Calculation Intervals in the hopping window that had a PM Metric Value that is equal to or greater than the TCA Performance Threshold Value (shown as 2 PMCIs in blue text in Figure 16). Moving forward 10 more PM Metric Calculation Intervals to interval 43, a TCA is not generated since zero of the PM Metric Calculation Intervals in the hopping window had a PM Metric Value is equal to or greater than the TCA Performance Threshold Value. At the completion of a hopping window without the occurrence of any TCAs, the damping algorithm resets to the start of the algorithm. If a future PM Metric Calculation Interval has a PM Metric Value equal to or greater than the TCA Performance Threshold Value, a TCA is generated at the completion of the PM Metric Calculation Interval, and the Damping Factor is activated. This behavior repeats as long as the Stateless TCA is enabled which is shown by the TCA at the completion of PM Metric Calculation Interval 45 which activates the Damping Factor again.

7.4.3 Threshold Crossing Alert Requirements

This section defines the requirements for the use of TCAs.

7.4.3.1 Common TCA Requirements

The requirements in this section apply to Stateful and Stateless TCA Reporting implementations.

[D3] An SD-WAN PM Implementation **SHOULD** provide the ability to generate TCAs based on the TCA Function Parameters used.

[CR1]<[D3] An implementation supporting TCAs **MUST** support at least one of *Stateful* or *Stateless* TCA Reporting.

[CR2]<[D3] An implementation supporting TCAs **MUST** support TCA Functions with different TCA Function Parameter values for each Monitored Entity and each PM Metric.

[CD1]<[D3] An implementation supporting TCAs **SHOULD** support multiple TCA Functions with different TCA Function Parameter values for a given Monitored Entity and PM Metric.

[CR3]<[D3] If a TCA is generated, then it **MUST** be generated at the end of the PM Metric Calculation Interval that triggered it.

The SD-WAN SP can report the occurrence of a TCA to the Subscriber.

[CR4]<[D3] If TCAs are reported to the Subscriber by the SD-WAN SP, they **MUST** include the information contained in either Table 11 or Table 13, and if applicable Table 14.

Note: the method used to display TCAs to the Subscriber is beyond the scope of this document.

7.4.3.2 Stateful TCA Requirements

The requirements in this section apply to Stateful TCA implementations.

[CR5]<[D3] An implementation supporting Stateful TCA Reporting **MUST** support SET-TCA functionality defined in section 7.4.1.

[CR6]<[D3] An implementation supporting Stateful TCA Reporting **MUST** support CLEAR-TCA functionality defined in section 7.4.1.

[CR7]<[D3] An implementation supporting Stateful TCA Reporting **MUST** support the value of the SET-TCA Window Thresholds being any integer value within a range of 1-300.

[CR8]<[D3] An implementation supporting Stateful TCA Reporting **MUST** support the value of the CLEAR-TCA Window Thresholds being any integer value within a range of 1-300.

[CR9]<[D3] An implementation supporting Stateful TCA Reporting **MUST** support the value of the TCA Window Size being any integer value within a range of 1-300.

[CR10]<[D3] An implementation supporting Stateful TCA Reporting **MUST** support TCA parameters for the PM Metric shown in Table 9.

PM Metric	Performance Threshold Value	SET-TCA Window Threshold Value	CLEAR-TCA Window Threshold Value	TCA Window Size Value
One-way Mean Packet Delay	PT_{MPD}	SWT_{MPD}	CWT_{MPD}	TWS_{MPD}

PM Metric	Performance Threshold Value	SET-TCA Window Threshold Value	CLEAR-TCA Window Threshold Value	TCA Window Size Value
One-way Mean Inter-Packet Delay Variation	PT_{IPDV}	SWT_{IPDV}	CWT_{IPDV}	TWS_{IPDV}
One-way Packet Loss Ratio	PT_{PLR}	SWT_{PLR}	CWT_{PLR}	TWS_{PLR}

Table 9 – Stateful TCA Reporting PM Metric Parameters

[CR11]<[D3] An implementation supporting Stateful TCA Reporting **MUST** support SET-TCA and CLEAR-TCAs when the conditions occur as shown in Table 10 and Figure 10.

PMV is used to abbreviate PM Metric Value in the following table.

PM Metric	TCA SET Criteria	TCA CLEAR Criteria
One-way Mean Packet Delay	<i>When $PMV_{MPD} \geq PT_{MPD}$</i>	<i>When $PMV_{MPD} < PT_{MPD}$</i>
One-way Mean Inter-Packet Delay Variation	<i>When $PMV_{IPDV} \geq PT_{IPDV}$</i>	<i>When $PMV_{IPDV} < PT_{IPDV}$</i>
Packet Loss Ratio	<i>When $PMV_{PLR} \geq PT_{PLR}$</i>	<i>When $PMV_{PLR} < PT_{PLR}$</i>

Table 10 – Stateful TCA Reporting SET & CLEAR Criteria

[CR12]<[D3] For an implementation supporting Stateful TCA Reporting, the following inequalities **MUST** hold:

- $SWT_{MPD} + CWT_{MPD} > TWS_{MPD}$
- $SWT_{IPDV} + CWT_{IPDV} > TWS_{IPDV}$
- $SWT_{PLR} + CWT_{PLR} > TWS_{PLR}$

[CR13]<[D3] An implementation supporting Stateful TCA Reporting **MUST** generate a SET-TCA notification message and set the internal state of the TCA Function to ‘set’ when the TCA Function is in the ‘clear’ state and the criteria for SET-TCA defined in Table 10 are met.

[CR14]<[D3] An implementation supporting Stateful TCA Reporting **MUST** generate a CLEAR-TCA notification message and set the internal state of the TCA Function to ‘clear’ when the TCA Function is in the ‘set’ state and the criteria for CLEAR-TCA defined in Table 10 are met.

[CR15]<[D3] An implementation supporting Stateful TCA Reporting **MUST** include the attributes shown in Table 11 in the SET-TCA or CLEAR-TCA notification message.

Field Name	Field Value	Field Description
Date and Time	Date/time in UTC	Time of the event, in UTC. For Stateful SET-TCA and CLEAR-TCA this is the time of the completion of the PM Metric Calculation Interval for which the PM Metric Value triggered the TCA to be generated.
Performance Monitoring Metric Name	One of One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, or One-way Packet Loss Ratio as a String	Human readable text for the Performance Monitoring Metric for which the TCA Function was configured, i.e., one of those listed in Table 9.
TCA Performance Threshold Value	Numeric value as an Integer ≥ 1	The TCA Performance Threshold Value for the Performance Monitoring Metric. Milliseconds are used for delay and percent us used for loss.
SET-TCA Window Threshold Value	Numeric value as an Integer ≥ 1	The value of the SET-TCA Window Threshold. Only used for SET-TCA notification messages.
CLEAR-TCA Window Threshold Value	Numeric value as an Integer ≥ 1	The value of the CLEAR-TCA Window Threshold. Only used for CLEAR-TCA notification messages.
TCA Window Size Value	Numeric value as an Integer ≥ 1	The number of PM Metric Calculation Intervals included in the sliding window for the SET-TCA or CLEAR-TCA process.

Field Name	Field Value	Field Description
PM Metric Value	List of Numeric value for each PM Metric Calculation Interval as an Integer	Units are delay or percentage
TCA Type	STATEFUL-SET, or STATEFUL-CLEAR as a String	The type of TCA, i.e., STATEFUL-SET or STATEFUL-CLEAR
Severity Level	CRITICAL, MAJOR, MINOR, WARNING, or CLEARED as a String	CRITICAL, MAJOR, MINOR, or WARNING apply to STATEFUL-SET, CLEARED applies to STATEFUL-CLEAR. These are agreed to by the Subscriber and SP.

Table 11 – Stateful TCA Notification Message Fields

Note: CLEARED is included in the Severity Level to align with TCA implementations where the STATEFUL-SET is a WARNING and the STATEFUL-CLEAR is shown as CLEARED.

7.4.3.3 Stateless TCA Reporting Requirements

The requirements in this section apply to Stateless TCA implementations.

[CR16]<[D3] An implementation supporting Stateless TCA Reporting **MUST** support the TCA functionality defined in section 7.4.2.

[CR17]<[D3] An implementation supporting Stateless TCA Reporting **MUST** trigger TCAs for the PM Metrics shown in Table 12 and Figure 13.

Note: PMV shown in Table 12 is the acronym for PM Metric Value.

PM Metric	TCA Performance Threshold Value	TCA Triggered
One-way Mean Packet Delay	PT_{MPD}	When $PMV_{MPD} \geq PT_{MPD}$
One-way Mean Inter-Packet Delay Variation	PT_{IPDV}	When $PMV_{IPDV} \geq PT_{IPDV}$
One-way Packet Loss Ratio	PT_{PLR}	When $PMV_{PLR} \geq PT_{PLR}$

Table 12 – Stateless TCA Reporting PM Metric Conditions

[CD2]<[D3] An implementation supporting Stateless TCA Reporting **SHOULD** support the Damping Factor.

[CR18]<[D3] An implementation supporting Stateless TCA Reporting with the Damping Factor **MUST** trigger TCAs for the PM Metrics shown in Table 12 and Figure 14.

[CR19]<[D3] An implementation supporting Stateless TCA Reporting with the Damping Factor **MUST** support the Damping Factor integer values between 2-300 PM Metric Calculation Intervals.

[CR20]<[D3] An implementation of Stateless TCA Reporting that includes the Damping Factor **MUST** generate a TCA at the end of a PM Metric Calculation Interval if the PM Metric Value for that PM Metric Calculation Interval is greater than or equal to the TCA Performance Threshold and no TCA has been generated after any of the $d-1$ preceding PM Metric Calculation Intervals, where d is the value of the Damping Factor.

[CR21]<[D3] An implementation of Stateless TCA Reporting that includes the Damping Factor **MUST NOT** generate a TCA at the end of a PM Metric Calculation Interval if a TCA has been generated after any of the $d-1$ preceding PM Metric Calculation Intervals, where d is the value of the Damping Factor.

[CR22]<[D3] An implementation of Stateless TCA Reporting that includes the Damping Factor **MUST** generate a TCA at the end of a PM Metric Calculation Interval if a TCA was generated at the end of the PM Metric Calculation Interval that precedes this one by d PM Metric Calculation Intervals, and for at least one of the $d-1$ intervening PM Metric Calculation Intervals, the PM Metric Value was greater than or equal to the TCA Performance Threshold, where d is the value of the Damping Factor.

[CR23]<[D3] An implementation of Stateless TCA Reporting that includes the Damping Factor **MUST NOT** generate a TCA at the end of a PM Metric Calculation Interval if the PM Metric Value was less than the TCA Performance Threshold for this PM Metric Calculation Interval and for all of the preceding $d-1$ PM Metric Calculation Intervals, where d is the value of the Damping Factor.

[CR24]<[D3] An implementation of Stateless TCA Reporting that does not include the Damping Factor **MUST** generate a TCA at the end of a PM Metric Calculation Interval if, and only if, the PM Metric Value for that PM Metric Calculation Interval is greater than or equal to the TCA Performance Threshold.

[CD3]<[D3] For an implementation of Stateless TCAs that includes the Damping Factor, when a TCA is generated at the end of a PM Metric Calculation Interval, it **SHOULD** include the number of PM Metric Calculation Intervals, within the sequence of d PM Metric Calculation Intervals ending with the one at the end of which the TCA was generated, in which the PM Metric Value was greater than or equal to the TCA Performance Threshold, where d is the value of the Damping Factor.

[CR25]<[D3] An implementation of Stateless TCA Reporting **MUST** include the information shown in Table 13 in the TCA.

Field Name	Field Value	Field Description
Date and Time	Date/time in UTC	Time of the event, in UTC. This is the time of the end of the PM Metric Calculation Interval for which the TCA is generated.
Performance Monitoring Metric Name	One of One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, or One-way Packet Loss Ratio as a String	Human readable text for Performance Monitoring Metric for which the TCA Function was configured, i.e., one of those listed in Table 12.
TCA Performance Threshold Value	Numeric value as an Integer ≥ 1	The TCA Performance Threshold Value
Performance Monitoring Metric Value	Numeric value as an Integer ≥ 1	The PM Metric Value for the PM Metric Calculation
TCA Type	STATELESS as a String	The type of TCA
Severity Level	One of CRITICAL, MAJOR, MINOR, WARNING as a String	CRITICAL, MAJOR, MINOR, or WARNING. Agreed to by Subscriber and SP.

Table 13 – Stateless TCA Reporting Notification Message Fields

[CR26]<[D3] If an implementation of Stateless TCA Reporting includes the Damping Factor, the Damping Factor shown in Table 14 **MUST** be appended to the TCA notification message.

[CR27]<[D3] If an implementation of Stateless TCA Reporting includes reporting the number of PM Metric Calculation Intervals per [CD3]<, that had a PM Metric Value that is equal to or greater than the TCA Threshold Value, the Number of PM Metric Calculation Intervals shown in Table 14 **MUST** be appended to the TCA Notification.

Field Name	Field Value	Field Format	Field Description
Damping Factor	Numeric value	Integer	The value that identifies the number of PM Metric Calculation Intervals included in the Damping Factor process.
Number of PM Metric Calculation Intervals	Numeric value	Integer	The number of PM Metric Calculation Intervals in the window defined by the Damping Factor in which the PM Metric Value \geq the TCA Performance Threshold Value

Table 14 – Damping Factor TCA Notification Message Field

As discussed in section 7.4, the Damping Factor can be used to reduce the number of TCAs generated by an implementation supporting a Stateless TCA Reporting. The relatively short duration of the PM Metric Calculation Interval can cause many TCAs to be generated over a short time period. The use of the Damping Factor mitigates this issue. There may be questions about how many of the PM Metric Calculation Intervals within the hopping window actually met the criteria to generate a TCA to understand the severity of the degradation or fault. For this reason, the Number of PM Metric Calculation Intervals attribute identifies the number of PM Metric Calculation Intervals that met the criteria for a TCA. An implementation of the Damping Factor without this optional capability may cause an SD-WAN SP to perform additional research to determine whether the number of PM Metric Calculation Intervals within the hopping window that met the criteria to generate a TCA.

8 Service Readiness Testing for an SD-WAN Service

Service Readiness Testing (SRT) is the process of testing SD-WAN Service to ensure it is ready for the Subscriber to begin using the service. It is performed to ensure that connectivity across UCSs exists between SD-WAN Edges. SRT for SD-WAN Service includes reporting UCS Service Attributes, SRT Parameters (defined in section 8.5), and SRT Results (defined in section 8.5) to the Subscriber. SRT verifies the connectivity of each ordered pair of UCS End Points within the service agreed to be Service Readiness Tested by the Subscriber and SD-WAN SP.

As a prerequisite, the SRT cannot begin until the SD-WAN Edges that are part of the Subscriber's SD-WAN Service are installed, have connectivity to the SD-WAN SP's SD-WAN Controller/Orchestrator, are configured with the basic SD-WAN Edge configuration used by the SD-WAN SP, and UCSs are connected to the SD-WAN Edges via UCS UNIs as appropriate for the SD-WAN Service.

A test methodology is defined for SRT for ordered pairs of UCS End Points within this document. This methodology provides a step-by-step process for performing a specific test or measurement. It also includes the SRT Parameters used for the test methodology.

The remainder of this section contains the following:

- A discussion of SRT use cases (section 8.1)
- A discussion of SRT terms and components (section 8.2)
- A description of Service Readiness Measurement Points (SRMPs) (section 8.3)
- A description of where SRMPs are located (section 8.3)
- Requirements for devices and applications including SRMPs (section 8.3)
- Tables that define which UCS attributes are tested, and which are reported (section 8.4)
- SRT for verifying connectivity of ordered pairs of UCS End Points (section 8.6)
- Test result reporting (section 8.7)

8.1 Verifying ordered pair of UCS End Points Connectivity

This use case describes the verification of ordered pair of UCS End Points connectivity. SRT Parameters and SRT Results are defined in section 8.5. As a new SD-WAN Service is being installed and before it is activated for the Subscriber to use, the SD-WAN SP verifies that there is connectivity between the agreed-on pairs of SD-WAN Edges used to implement the SD-WAN Service. This is accomplished by performing SRT on the agreed on ordered pairs of UCS End Points.

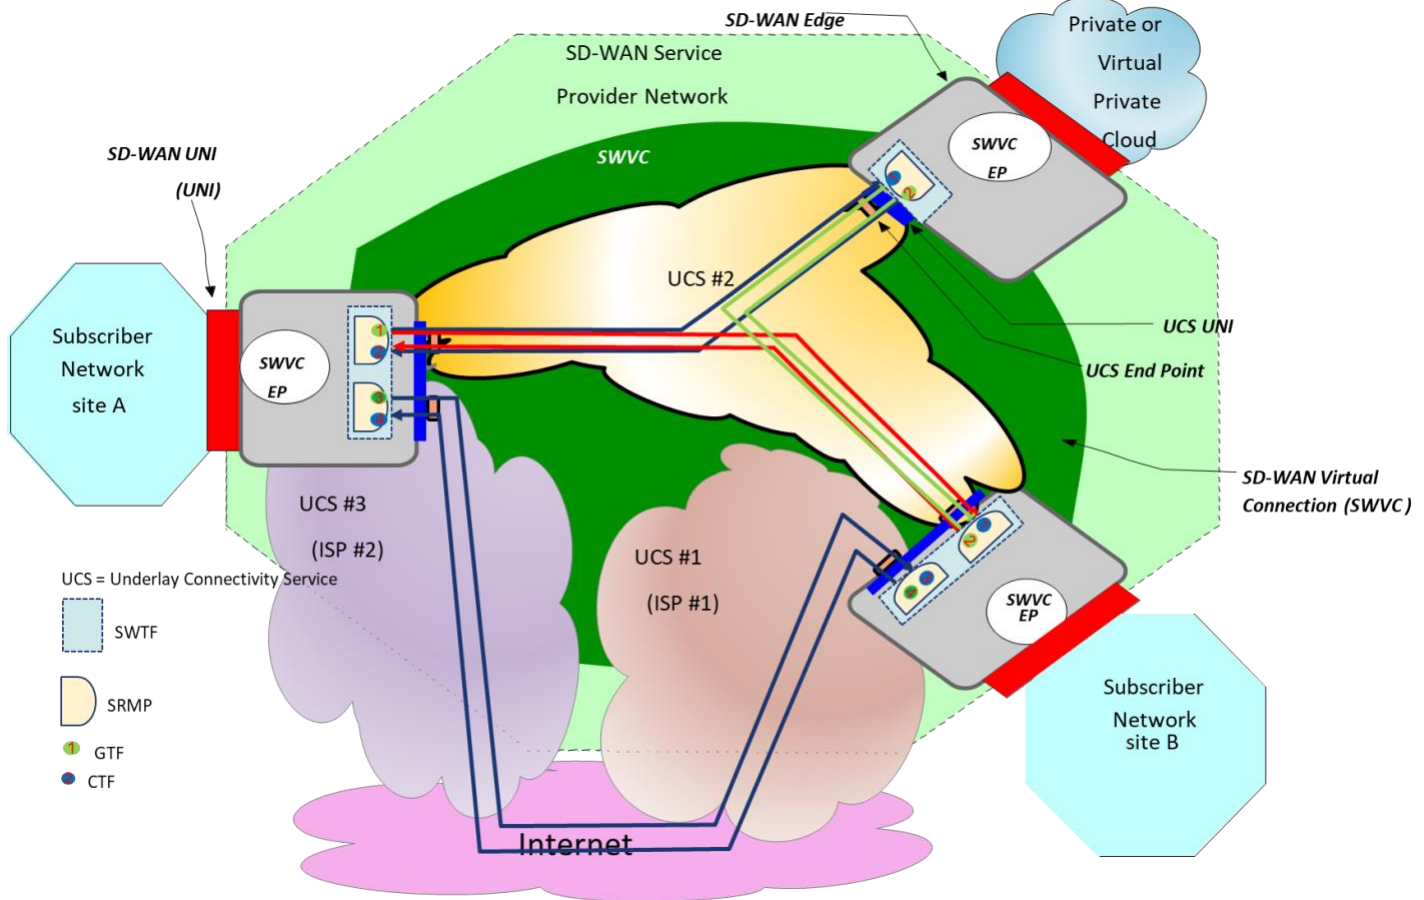


Figure 17 – SRT Order Pairs of UCS End Points Connectivity Verification Example

Figure 17 shows the use case for the verification of connectivity between SD-WAN Edges. SRMPs (as described in section 8.2) are used at each of the UCS End Points in the ordered pair and tests are performed on each ordered pair of UCS End Points. When a new SD-WAN Service is being activated, tests are performed on all agreed on ordered pairs of UCS End Points.

When a new SD-WAN Edge is added to an existing SWVC, the ordered pairs of UCS End Points between that SD-WAN Edge and existing SD-WAN Edges it is connected to may be tested. It is suggested that if this testing is performed, downtime with the Subscriber be arranged to avoid disrupting any Subscriber traffic that is sharing the same UCS.

8.2 Service Readiness Testing Terms and Components

This section describes terms and components used to perform SRT. SRT is performed using at least two Service Readiness Measurement Points (SRMPs). The SRMP is a logical point inside an SD-WAN Edge. The SRMP also contains both a Generator Test Function (GTF) and a Collector Test Function (CTF). A GTF generates IP Test Packets used for test measurements. A CTF either counts and discards IP Test Packets coming from a GTF or counts and processes IP Test Packets from a GTF. When testing with Unicast IP Test Packets, a GTF is paired with a CTF so that the IP Test Packets generated by the GTF are collected by a particular CTF.

An SRT Methodology is defined for connectivity testing and verification of ordered pairs of UCS End Points. The SRT Methodology identifies the test name, test objective, test procedure, variables used in the methodology, results, and remarks. SRT Methodologies are specified in section 8.6.

8.3 Service Readiness Measurement Point Locations

The logical location of SRMPs within the network is shown in this section. The following figures show the location of SRMPs in relationship to processing functions within the SD-WAN Edge. The SRMPs are located so that IP Test Packets pass over the UCS which connects the ordered pair of UCS End Points between two SD-WAN Edges but are not processed by functions associated with the SWVC End Point. How these functions are implemented is outside the scope of this document.

The tool used to generate and receive packets is beyond the scope of this document.

- [R29] The SRMP **MUST** be located so that IP Test Packets generated by the GTF are inserted at the UCS UNI which terminates the UCS under test.
- [R30] The SRMP **MUST** be located so that IP Test Packets collected by the CTF are received at the UCS UNI which terminates the UCS under test.

The SRMP used for SRT connects to a UCS UNI. It is located so that IP Test Packets generated by the GTF are inserted on an UCS UNI without being processed by the functions associated with the SWVC End Point. IP Test Packets collected by the CTF are received directly from the UCS UNI and are not processed by the functions associated with the SWVC End Point. This is shown in Figure 18.

If any TVCs exist at the time that SRT is performed, the SRMP is located so that it connects to the UCS UNI and not to the TVC(s).

- [R31] IP Test Packets **MUST** be formatted so that they are mapped to the appropriate UCS End Point within a UCS UNI.

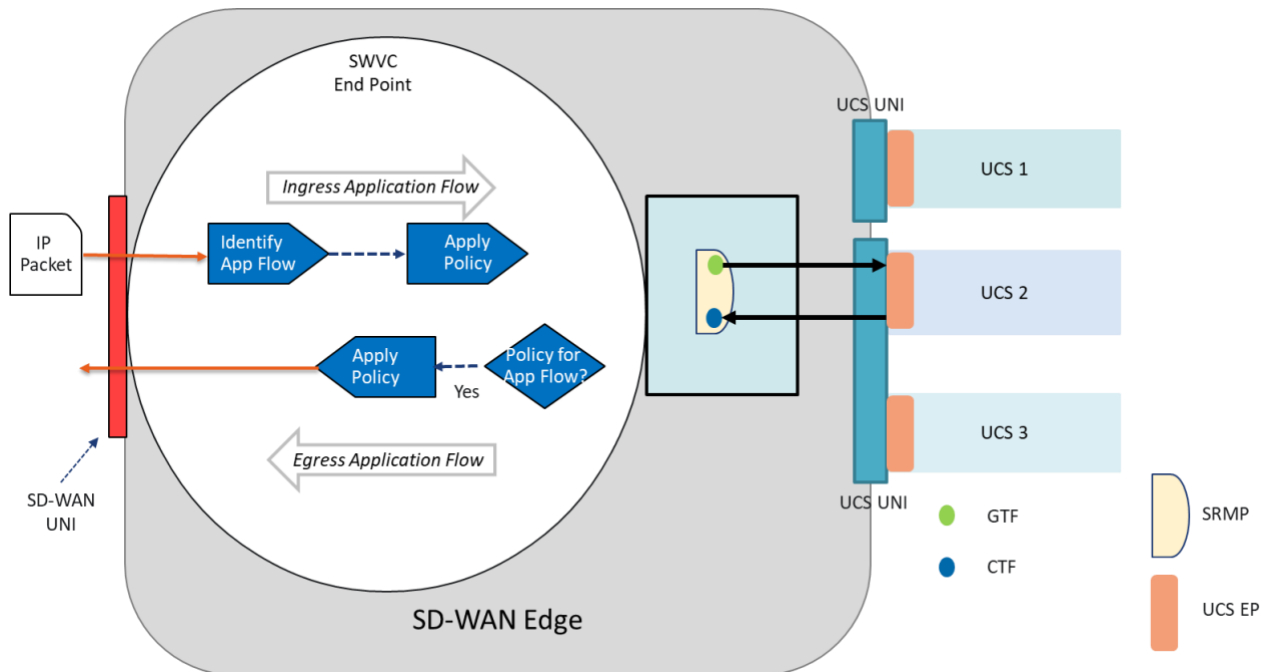


Figure 18 – SRMP Location of UCS End Points Connectivity Testing

[R32] An SRMP **MUST** contain a Generator Test Function (GTF) and a Collector Test Function (CTF).

8.4 UCS Service Attribute Reporting

This section of the document details the UCS, UCS UNI, and UCS End Point Service Attributes that are included in the SRT report as a part of the SRT Process.

[R33] The value of all UCS Service Attributes, UCS UNI Service Attributes and UCS End Point Service Attributes defined in MEF 70.1 [13] **MUST** be reported to the Subscriber as part of the SRT methodology, for a new SD-WAN Service or when a new SD-WAN UNI or a new UCS UNI is added to an existing SD-WAN Service.

Note: UCS Service Attributes are not verified by SRT, the values of the attributes are reported.

8.5 Service Readiness Testing Parameters and Results

SRT verifies and reports the results of the UCS End Point ordered pair Connectivity tests. A prerequisite to this testing is verifying that the SD-WAN Controller/Orchestrator can communicate with SD-WAN Edges that are a part of the SWVC.

8.5.1 SD-WAN Edge to SD-WAN Controller/Orchestrator Communication

As a prerequisite to SRT, the verification of communication between the SD-WAN Controller/Orchestrator and SD-WAN Edges is performed. The steps used to verify this

communication are beyond the scope of this document. The SRT report does not indicate if this verification was performed or if it passed or failed.

8.5.2 Ordered Pairs of UCS End Points SRT

As discussed previously in section 8, the SD-WAN SP and Subscriber agree on the ordered pairs of UCS End Points that are included in SRT. Verifying connectivity of the ordered pairs between SD-WAN Edges for a given SWVC is required for SRT. The ordered pairs of UCS End Points are either part of two different Internet Access UCSs or a non-Internet UCS. SRT verifies that IP Test Packets can be sent from the first UCS End Point in the ordered pair to the second UCS End Point. The Pass/Fail result of each tested ordered pair of UCS End Points is reported.

- [R34] When a new SWVC is being activated, the SD-WAN SP **MUST** verify connectivity between the agreed upon ordered pairs of UCS End Points.
- [R35] When a new SD-WAN Edge is added to an existing SWVC, the SD-WAN SP **MUST** verify connectivity between the agreed upon ordered pairs of UCS End Points.
- [R36] When the SD-WAN SP is verifying connectivity, they **MUST** use the Test Methodology defined in Table 15.

8.6 Test Methodology

This section contains the Test Methodology for the verification of the service readiness including the methodology for connectivity verification of ordered pairs of UCS End Points.

Figure 19 shows the high level UCS End Point ordered pair Connectivity SRT methodology.

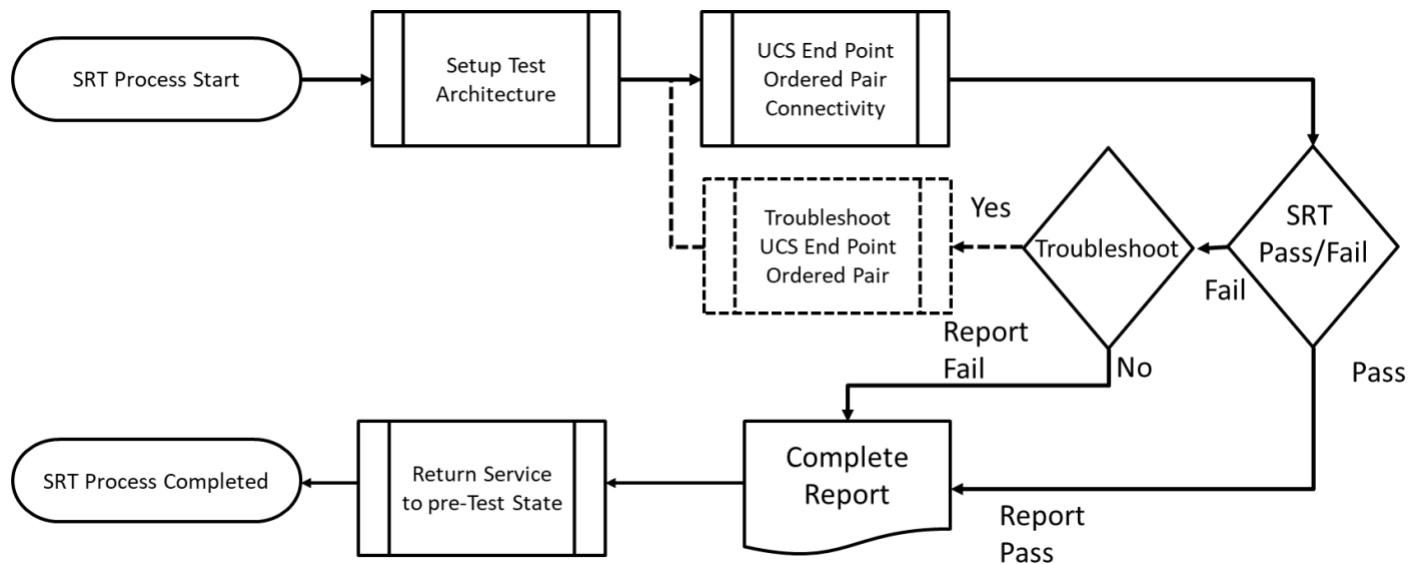


Figure 19 – UCS End Point Ordered Pair Single Direction Connectivity SRT Methodology

SRT is repeated for each ordered pair that is agreed on be included in SRT by the Subscriber and SD-WAN SP.

The first step of the methodology is to create the required SRMPs at the appropriate SD-WAN Edges and to connect the SRMP to the correct UCS UNI.

The second step in the methodology is to verify connectivity between the end point pair as described in section 8.6.1.

The third step in the methodology is to report the UCS Service Attributes (see Complete Report step in Figure 19), and SRT Results.

The fourth and final step in the methodology is to disconnect the SRMP from the UCS UNI and to restore the UCS End Point to its normal configuration.

Troubleshooting of a test failure is optional and depends on factors such as who provided the UCS and who is responsible for UCS management. If troubleshooting of the failure is performed, the UCS End Point ordered pair test is repeated when the trouble has been resolved. If troubleshooting is not performed, the failure is reported.

8.6.1 UCS End Point Ordered Pair Connectivity Testing

Connectivity between ordered pairs of UCS End Points is tested for the ordered pairs that are agreed to be tested by the SD-WAN SP and the Subscriber. The verification of connectivity uses the methodology defined in Table 15.

[R37] Results for each ordered pair of UCS End Points tested in the SRT Test Methodology **MUST** be reported as *pass* or *fail*.

Service Readiness Test Methodology	
Test Name	UCS End Point ordered pair Connectivity
Test Objective	Verify that there is connectivity between SD-WAN Edges per ordered pair of UCS End Points
Test Procedure	<ul style="list-style-type: none"> •For this Test Methodology SRMP₁ and SRMP₂ are placed as shown in Figure 18.. •SRMP₁ offers a number of IP Test Packets with the DA for reaching SRMP₂ so that the IP Test Packets are injected into the UCS End Point at location 1 at an interval where the IP Test Packets are equally distributed over time T_{SC}. The number of IP Test Packets offered and the value of time T_{SC} are agreed to by the SP and Subscriber. •SRMP₂ counts the IP Test Packets received from SRMP₁ and either provides the count of received packets or subtracts the number of received packets from the number of offered packets and provides the number of lost packets. Note: If SRMP₂ does not know the number of offered IP Test Packets, then the calculation of lost packets is performed at some point that is aware of both the number of offered IP Test Packets and the number of received IP Test Packets. •The above is repeated for each ordered pair of UCS End Points in each direction agreed to be tested by the Subscriber and SD-WAN SP.
Parameters	Set of order pairs of UCS End Points, T_{SC} , number of IP Test Packets Offered, number of lost packets allowed per ordered pair of UCS End Points
Results	Pass = If the number of lost packets is less than or equal to the allowed value agreed to by Subscriber and SP Fail = If the number of lost packets is greater than the allowed value
Remarks	

Table 15 – UCS End Point Ordered Pair Connectivity Test Methodology

[R38] The SD-WAN SP and the Subscriber **MUST** agree on the parameters of Service Readiness Testing.

Note: To test the UCS End Point ordered pair in each direction, this methodology is repeated for the other direction between the SD-WAN Edges over the same UCS.

When the agreed-on set of order pairs of UCS End Points have been verified, the Subscriber can begin to use the service.

8.7 Test Record

After all tests have been completed an SRT record is created. The SRT record contains the attribute and test result information described in sections 8.4, 8.5, and 8.6. The results from the different tests are mapped into one SRT record for that service. The SRT record can be shared with the Subscriber and can be stored within SD-WAN SP management systems. The format of the SRT record is not mandated by this document.

9 References

- [1] IETF RFC 2119, *Key words for use in RFCs to Indicate Requirement Levels*, by S. Bradner, March 1997
- [2] IETF RFC 3339, *Date and Time on the Internet: Timestamps*, by G. Clyne, C. Newman, July 2002, Copyright (C) The Internet Society (2002). All Rights Reserved.
- [3] IETF RFC 8174, *Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words*, by B. Leiba, May 2017, Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.
- [4] ISO/IEC 7498-1:1994, *Information technology – Open System Interconnection – Basic Reference Model: Basic Model*, June 1996
- [5] ITU-T Recommendation X.733, *Information Technology – Open Systems Interconnection, System Management – Alarm Reporting Function*, 1992
- [6] MEF 10.4, *Subscriber Ethernet Service Attributes*, December 2018
- [7] MEF 35.1, *Service OAM Performance Monitoring Implementation Agreement*, May 2015
- [8] MEF 48.1, *Carrier Ethernet Service Activation Testing*, November 2019
- [9] MEF 55.1, *Lifecycle Service Orchestration (LSO): Reference Architecture and Framework*, January 2021
- [10] MEF 61.1, *IP Service Attributes*, May 2019
- [11] MEF 66, *Service OAM for IP Services*, July 2020
- [12] MEF 67, *Service Activation Testing for IP Services*, December 2020
- [13] MEF 70.1, *SD-WAN Service Attributes and Service Framework*, November 2021

Appendix A Service Readiness Testing vs Service Activation Testing

This document introduces the concept of Service Readiness Testing for SD-WAN Services. This is different from Service Activation Testing that is defined for other MEF services. Service Activation Testing is performed on the service and verifies that the service is operating as described by the agreed to Service Attributes. The Service Attributes that are reported and/or tested are included in the definition of Service Activation Testing.

SD-WAN Service is built to run on top of one or more UCSs and uses Policies to determine how an IP Packet received from a Subscriber is forwarded or discarded. A SD-WAN SP might maintain many Policies. Some might be default Policies used for all Subscribers and others might be developed to address a specific Subscriber's application. Verifying the correct operation of each Policy may be done in a lab environment or in conjunction with a specific Subscriber's applications.

Rather than performing Service Activation Testing on SD-WAN Service, Service Readiness Testing is performed. Service Readiness Testing verifies that the SD-WAN Service is ready for the Subscriber to use or the SD-WAN SP or Subscriber to implement the Subscriber-specified Policies. SRT does not verify the operation of the Policies or SD-WAN Service Attributes. Instead, it is focused on determining that each ordered pair of UCS End Points provides connectivity between the appropriate SD-WAN Edges. It does report the UCS Service Attributes.

The results of the SRT may be presented to the Subscriber as an SRT report. The contents of the SRT Report are defined in this document. The format of the SRT report is beyond the scope of this document.

The scope of testing within this document is limited to Service Readiness Testing. Any additional testing that the SD-WAN SP and Subscriber agree to perform is beyond the scope of this document.