



MEF 105 Draft (R5)

Performance Monitoring and Service Readiness Testing for SD-WAN

January 2023

**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 2023. 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.....	2
3	Terminology and Abbreviations.....	3
4	Compliance Levels	8
5	Conventions	9
5.1	Numerical Conventions	9
5.2	Diagram Conventions	9
6	Introduction.....	0
7	SD-WAN Performance Monitoring.....	3
7.1	Performance Monitoring Framework	3
7.2	Performance Monitoring Metrics	4
7.2.1	Application Monitoring Flow and TVC Performance Monitoring.....	4
7.2.2	PM Definitions.....	17
7.2.2.1	<i>PM Metric List</i>	<i>17</i>
7.2.2.1.1	PM Metrics for Application Monitoring Flows	18
7.2.2.1.2	PM Metrics for Application Flows	19
7.2.2.1.3	PM Metrics for TVCs	20
7.2.2.2	<i>PM Metric Calculation Intervals</i>	<i>21</i>
7.3	PM Metric Calculation and Reporting.....	24
7.3.1	PM Metric Calculation Data Set	28
7.3.2	PM Metric Report	30
7.4	Threshold Crossing Alerts	31
7.4.1	Stateful Threshold Crossing Alert Reporting	32
7.4.2	Stateless Threshold Crossing Alert Reporting.....	35
7.4.3	Threshold Crossing Alert Requirements.....	38
7.4.3.1	<i>Common TCA Requirements</i>	<i>38</i>
7.4.3.2	<i>Stateful TCA Requirements</i>	<i>39</i>
7.4.3.3	<i>Stateless TCA Reporting Requirements.....</i>	<i>41</i>
8	Service Readiness Testing for an SD-WAN Service	45
8.1	Verifying ordered pair of UCS End Points Connectivity Use Case	45
8.2	Service Readiness Testing Terms and Components	46
8.3	Service Readiness Measurement Point Locations	47
8.4	UCS Service Attribute Reporting	48
8.5	Service Readiness Testing Parameters and Results	48
8.5.1	Ordered Pairs of UCS End Points SRT.....	48
8.6	Test Methodology	49
8.6.1	UCS End Point Ordered Pair Connectivity Testing.....	50
8.7	Test Record.....	52
9	References	53
Appendix A	Service Readiness Testing vs Service Activation Testing	54

List of Figures

Figure 1 – Diagram Conventions	0
Figure 2 – AFMP Locations	5
Figure 3 – Measured Information Rate Byte Counter Locations.....	6
Figure 4 – Ingress AFMP to Egress AFMP ordered pair Example	7
Figure 5 – TVC MP Locations.....	15
Figure 6 – Ingress TVC MP to Egress TVC MP ordered pair Example.....	16
Figure 7 – PM Metric Calculation Interval Concept	22
Figure 8 – PM Metric Calculation Durations	23
Figure 9 – PM Metric Calculation and Reporting Example	24
Figure 10 – PM Metric Calculation Interval Duration Equals PM Metric Report Interval Duration	26
Figure 11 – PM Metric Report Interval Greater Than PM Metric Calculation Interval Duration	27
Figure 12 – Different PM Metric Calculation Interval Durations in a PM Metric Report	28
Figure 13 – Flowchart for Stateful TCA.....	33
Figure 14 – Stateful TCA.....	34
Figure 15 – Stateful TCA without Consecutive PM Metric Calculation Intervals \geq TCA Performance Threshold Value Example	35
Figure 16 – Flowchart for Stateless TCA w/o Damping Factor	36
Figure 17 – Flowchart for Stateless TCA w/Damping Factor = <i>enabled</i>	36
Figure 18 – Stateless TCA with Damping Factor = <i>disabled</i> Example	37
Figure 19 – Stateless TCA with Damping Factor Example.....	37
Figure 20 – SRT Order Pairs of UCS End Points Connectivity Verification Use Case Example	46
Figure 21 – SRMP Location of UCS End Points Connectivity Testing	48
Figure 22 – UCS End Point Ordered Pairs Connectivity SRT Methodology	50

List of Tables

Table 1 - Terminology	6
Table 2 - Abbreviations	7
Table 3 – Application Flow Specification, Zone name and PM Metric Agreement Example	8
Table 4 – Monitored SD-WAN UNI ordered pair Agreement Example	9
Table 5 – Monitored SD-WAN UNI Agreement Example	9
Table 6 – Application Flow Specification, Zone name, ordered pair, and PM Metric Example .	11
Table 7 – Application Flow Specification Zone name, UNI Identifier, and PM Metric Overview Example	13
Table 8 – PM Metric Calculation Data Set Attributes	29
Table 9 – PM Metric Report Attributes	30
Table 10 – Stateful TCA Reporting PM Metric Parameters	40
Table 11 – Stateful TCA Notification Message Fields	41
Table 12 – Stateless TCA Reporting Notification Message Fields	43
Table 13 – Damping Factor TCA Notification Message Field	44
Table 14 – UCS End Point Ordered Pair Connectivity Test Methodology	51

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 [9], can use Performance Monitoring metrics to identify degradations or failures of TVCs and Application Flows. The Performance Metrics that are used to specify Application Flow performance objectives are defined in MEF 70.1 [9]. This document defines how the Performance Metrics are measured, calculated, and shared with Subscribers. 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 [9] and the Performance Monitoring metrics defined within this document is that the metrics are specified between different points. The MEF 70.1 [9] Performance Metrics are defined from SD-WAN UNI to SD-WAN UNI. The Performance Monitoring metrics defined within this document are defined between ordered pairs of 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 delivers to the SD-WAN Subscriber are detailed within this document.

In addition, Service Readiness Testing requirements for SD-WAN Services are defined within this document. The Service Readiness Testing requirements are focused on the IP level because MEF 70.1 [9] defines an SD-WAN Service as an IP service where ‘the basic unit of transport at the SD-WAN UNI is an IP Packet’. The SD-WAN Service Readiness Testing scope, in this standard, covers 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 [5], MEF 61.1 [6], MEF 66 [7], and MEF 70.1 [9] are included in this document by reference and are not repeated in the table below.

Term	Definition	Reference
Application Flow Measurement Point	A point where measurements of IP Packet performance are performed for Application Flows	This document
Application Monitoring Flow	This document defines an Application Monitoring Flow as a set of IP Packets 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.	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 [8]
Damping Factor	An integer > 0 that is used to suppress new TCAs. Applicable to Stateless TCA Reporting only.	This document
Egress Measured Information Rate	The Measured Information Rate at the Egress SD-WAN UNI	This document
Generator Test Function	A logical function for generating and transmitting IP Packets which can include test packets.	MEF 67 [8]
Information Rate	The rate, in bits per second, that IP Packets pass a measurement point.	This document
Ingress Measured Information Rate	The Measured Information Rate at the Ingress SD-WAN UNI	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

Term	Definition	Reference
Monitored Entity	<p>A Monitored Entity is one of the following:</p> <ol style="list-style-type: none"> 1. An Application Monitoring Flow flowing between a pair of UNIs, identified by the 4-tuple <i><Application Flow Specification, Zone name, Ingress UNI Identifier, Egress UNI Identifier></i> 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 3-tuple <i><Application Flow Specification, Zone name, UNI Identifier></i> - used for Ingress MIR and Egress MIR of an Application Flow 3. A TVC in a given direction, identified by the 3 tuples <i><Ingress UCS End Point Identifier, Egress UCS End Point Identifier, UCS CoS Name></i> - used for One-way Mean Packet Delay, One-way Inter-Packet Delay Variation and One-way Packet Loss Ratio of the TVC. 	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 (Application Flow Specification, Zone name, Ingress UNI Identifier, Egress UNI Identifier). For Application Flows, it is the 3-tuple (Application Flow Specification, Zone name, 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 [4])
Performance Monitoring Metric	A metric that is measured or calculated as a part of Performance Monitoring.	This document
PM Implementation	An implementation of the requirements for SD-WAN PM in this document.	This document
Performance Monitoring 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

Term	Definition	Reference
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 Duration	The duration of the time interval over which one or more PM Metrics are calculated.	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 Duration	The time interval over which one or more PM Metric values from one or more PM Metric Calculation Intervals are reported to the Subscriber by the SD-WAN Service Provider.	This document
SD-WAN Controller/Orchestrator	The entity that is responsible for managing/orchestrating the SD-WAN Service.	This document
SD-WAN Performance Monitoring	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 [5]
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
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 [4]

Term	Definition	Reference
Stateless TCA Reporting	A TCA reporting mechanism whereby TCAs are generated whenever an alertable condition is detected.	This document derived from MEF 35.1 [4]
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
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 [4]
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

Table 1 - Terminology

Term	Definition	Reference
AFMP	Application Flow Measurement Point	This document
CTF	Collector Test Function	MEF 67 [8]
EMIR	Egress Measured Information Rate	This document
GTF	Generator Test Function	MEF 67 [8]
IMIR	Ingress Measured Information Rate	This document
PM	Performance Monitoring	This document (derived from MEF 35.1 [4])
SD-WAN PM	SD-WAN Performance Monitoring	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
TCA	Threshold Crossing Alert	This document derived from MEF 35.1 [4]
TVC MP	Tunnel Virtual Connection Measurement Point	This document

Term	Definition	Reference
UTC	Coordinated Universal Time	IETF RFC 3339 [2]

Table 2 - Abbreviations

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 4.

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 several components that appear frequently. These components are represented in a standard way as described in Figure 1:

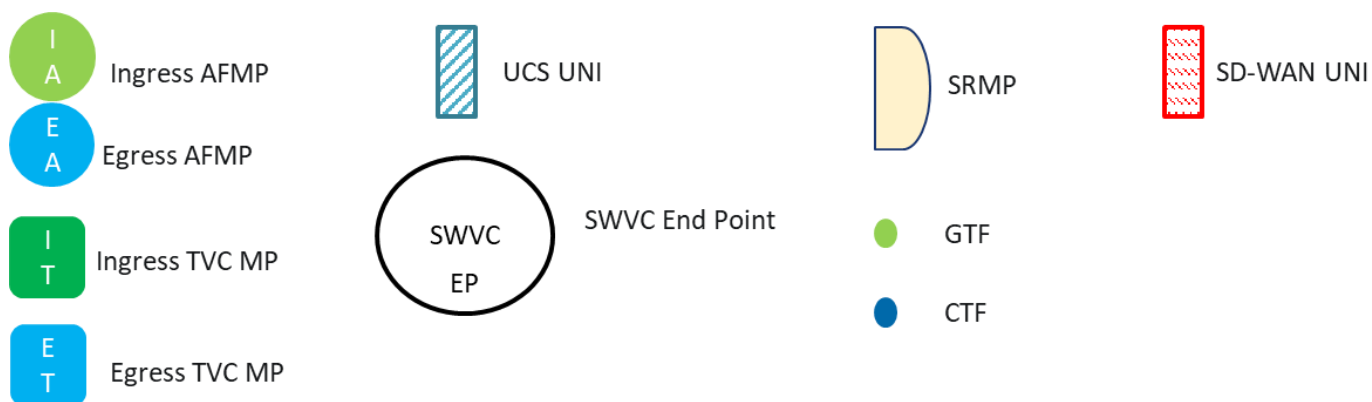


Figure 1 – Diagram Conventions

6 Introduction

For the SD-WAN Service standardized in MEF 70.1 [9], there is a need for 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 verifying 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 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.

The set of Application Monitoring Flows that are monitored is agreed to by the Subscriber and SP. Therefore, the ones that are monitored can be a strict subset of all Application Monitoring Flows.

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.

This document describes the requirements for the SD-WAN PM and SD-WAN SRT pertinent to an SD-WAN Service as described and defined in MEF 70.1 [9]. 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. These are:

- Performance Monitoring Metric Values per Application Monitoring Flow
- Performance Monitoring Metric Values per ordered pair of Tunnel Virtual Connection (TVC) end points
- 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 all UCS End Point pairs. This is because a TVC represents 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 measured over TVCs could be used to determine if a Path is meeting the PERFORMANCE Policy criterion specified for an Application Flow (see section 15 of MEF 70.1 [9] for details). Forwarding of IP Packets of a given Application Flow over the Path could be stopped if a constituent 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 [9]
- 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 [9]. Instead, it is focused on verifying that the UCSs providing connectivity between the appropriate SD-WAN Edges are working correctly. Appendix A contains a comparison of SRT and more traditional Service Activation Testing.

Areas that are not addressed within this document include the specific tools or implementations used to perform SD-WAN PM or SD-WAN SRT.

Note that there are several parameters values that are mandated to be agreed to by the Subscriber and SP. These parameters could be represented as Service Attributes but are not within the context of this standard.

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 Framework defined in section 7.1
- Performance Monitoring Metrics for Application Flow Specifications and TVCs defined in section 7.2
- Performance Monitoring Metric Calculation and Reporting defined in section 7.3
- 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. Application Monitoring Flows represent a given Application Flow Specification, Zone name combination. This is a 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 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 document.

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.

Note: 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

SD-WAN PM starts with measurement intervals that are significantly shorter (10 seconds or less) than the typical 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 and 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 specified 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.

Two types of Performance Monitoring are described in this document, Active Performance Monitoring and Passive Performance Monitoring. Active Performance Monitoring uses synthetic IP packet to perform measurements. Passive Performance Monitoring does not use synthetic IP Packets or modify Subscriber IP Packets to perform measurements.

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 (this is beyond the scope of this document).

7.2 Performance Monitoring Metrics

This section describes the measurement methods and PM Metrics that are defined for SD-WAN Service. The PM Metrics are:

- One-way Mean Packet Delay
- One-way Mean Inter-Packet Delay Variation
- One-way Packet Loss Ratio
- Measured Information Rate

The first three 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.

Note that the PM Metrics defined in this standard are different than the Performance Metrics (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 that are defined in MEF 70.1 [9] section 15.

The PM Metrics (One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, and One-way Packet Loss Ratio) for IP Packets flowing over either an Application Monitoring Flow or a TVC, are described in section 7.2.2.1. Measured Information Rate is also described in section 7.2.2.1.

7.2.1 Application Monitoring Flow and TVC Performance Monitoring

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 between Application Flow Measurement Points (AFMPs).

- An Ingress AFMP is located after where any Policies are applied and before the UCS UNI (see Figure 2).
- An 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 and 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 that do not use synthetic IP Packets or modify Subscriber IP Packets to perform measurements might result in PM Metrics that are reported as Undefined.

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

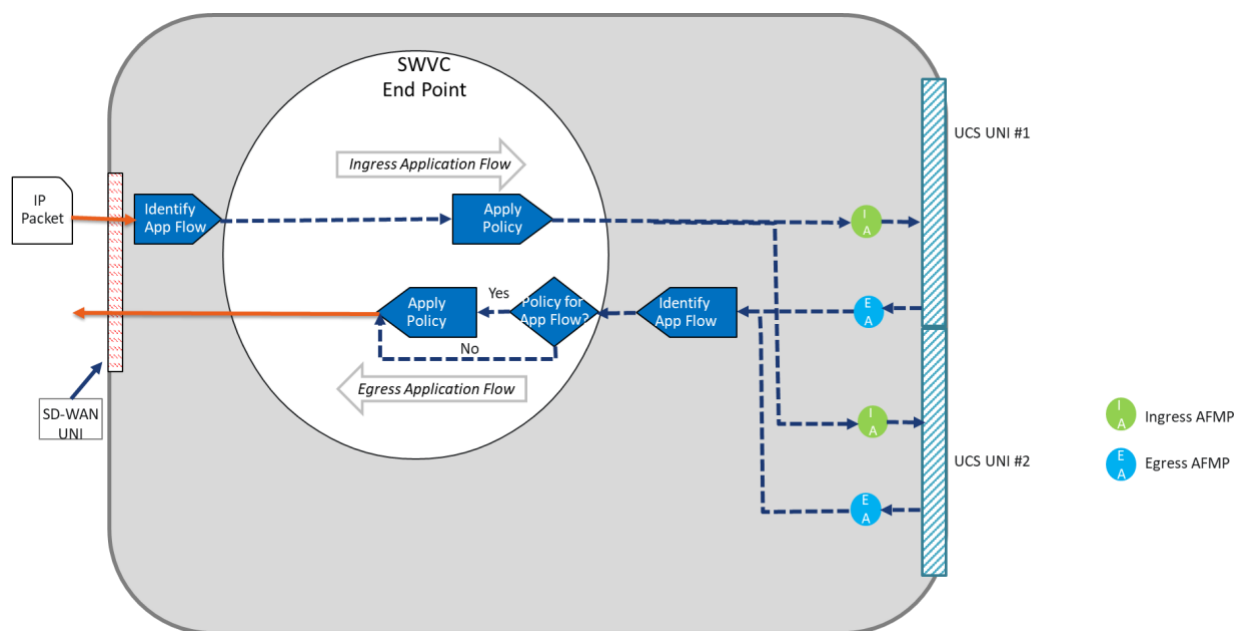


Figure 2 – AFMP Locations

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.

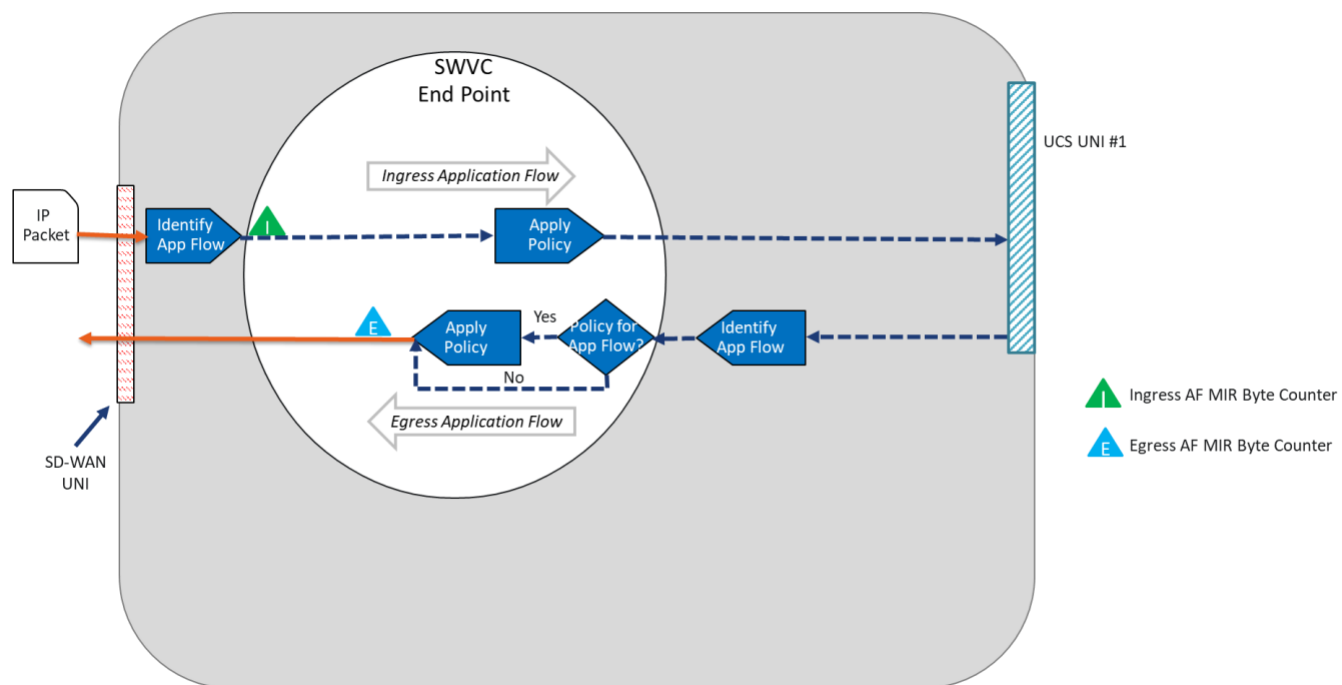


Figure 3 – Measured Information Rate Byte Counter Locations

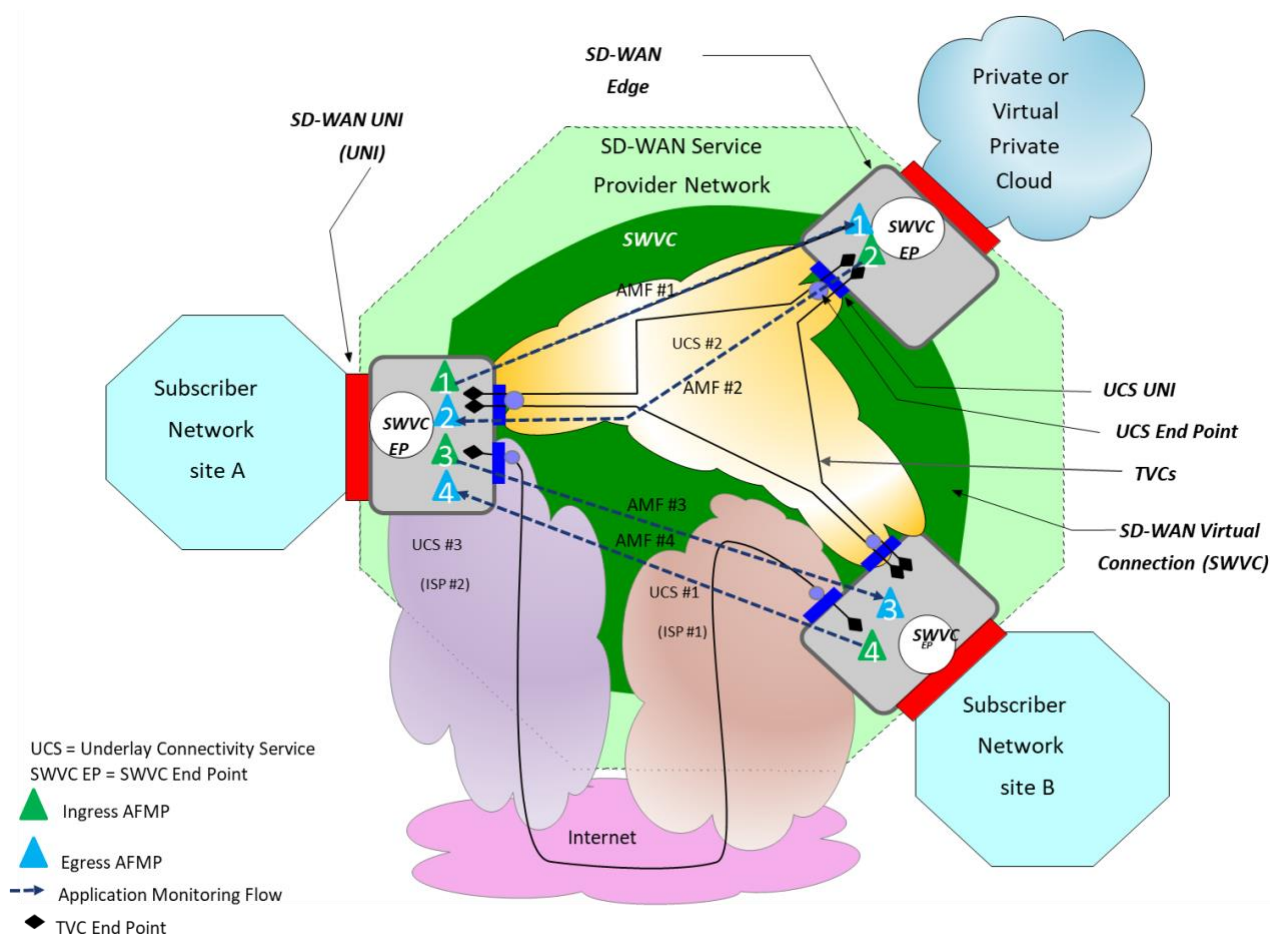


Figure 4 – Ingress AFMP to Egress AFMP ordered pair Example

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 (Application Flow Specification, Zone name x), AFMP ordered pair 2 monitors Application Monitoring Flow 2 (Application Flow Specification, Zone name y), AFMP ordered pair 3 monitors Application Monitoring Flow 3 (Application Flow Specification, Zone name z), and AFMP ordered pair 4 monitor Application Monitoring Flow 4 (Application Flow Specification, zone name A). For details on how one-way measurements are performed please see MEF 66 [7] section 8.3.

The following requirements apply to Application Monitoring Flows.

- [R1] The Subscriber and SP **MUST** agree on the set of Application Flow Specification, Zone name pairs to be monitored.
- [R2] For each Application Flow Specification, Zone name 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 3.

Application Flow Specification, Zone name	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 3 – Application Flow Specification, Zone name and PM Metric Agreement Example

- [R4] The Subscriber and SP **MUST** agree on a single set of SD-WAN UNI ordered pairs which will be monitored as described in [R6].

An example is shown in Table 4.

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

Table 4 – Monitored SD-WAN UNI ordered pair Agreement Example

Note that the combination of Application Flow Specification and Zone name (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 UNI ordered pairs which Ingress Measured Information Rate and Egress Measured Information Rate will be monitored per [R7].

The SP performs measurements between Ingress and Egress AFMPs for each ordered pair of UNIs as shown in Table 4. An example is shown in Table 5.

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

Table 5 – Monitored SD-WAN UNI Agreement Example

Note that the combination of Application Flow Specification and Zone Name (from [R1]) and a UNI (from [R5]) uniquely identifies an Ingress Application Flow and an Egress Application Flow.

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 6 shows an example of the PM Metrics that will be monitored for each Application Flow Specification, Zone name and SD-WAN UNI ordered pair (i.e., for each Application Monitoring Flow).

Application Flow Specification, Zone name	SD-WAN UNI ordered pair	PM Metric
Voice, Corporate	A-B	One-way Mean Packet Delay

Application Flow Specification, Zone name	SD-WAN UNI ordered pair	PM Metric
		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 One-way Packet Loss Ratio

Application Flow Specification, Zone name	SD-WAN UNI ordered pair	PM Metric
Video Streaming, Corporate	B-A	One-way Mean Packet Delay One-way Packet Loss Ratio
Video Streaming, Corporate	A-D	One-way Mean Packet Delay One-way Packet Loss Ratio
Video Streaming, Corporate	E-A	One-way Mean Packet Delay One-way Packet Loss Ratio
Video Streaming, Corporate	A-F	One-way Mean Packet Delay One-way Packet Loss Ratio
Video Streaming, Corporate	F-A	One-way Mean Packet Delay One-way Packet Loss Ratio

Table 6 – Application Flow Specification, Zone name, ordered pair, and PM Metric Example

Given the examples above, Table 7 shows an example of the PM Metrics that will be monitored for each Application Flow Specification, Zone name pair and a given SD-WAN UNI.

Application Flow Specification, Zone name	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

Application Flow Specification, Zone name	SD-WAN UNI	PM Metric
Voice, Corporate	D	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	E	Ingress Measured Information Rate Egress Measured Information Rate
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

Application Flow Specification, Zone name	SD-WAN UNI	PM Metric
		Egress Measured Information Rate
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 7 – Application Flow Specification Zone name, UNI Identifier, and PM Metric Overview Example

[R6] For a given pair of Application Flow Specification and Zone name, $\langle afs, z \rangle$, agreed per [R1], and a given ordered pair of UNIs $\langle u1, u2 \rangle$ agreed per [R4], the SP **MUST** monitor performance for the Application Monitoring Flow $\langle iaf, u2 \rangle$, where iaf is the Ingress Application Flow identified by the 3-tuple $\langle afs, z, u1 \rangle$, for each of the following 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 Zone name, $\langle afs, z \rangle$, agreed per [R1], and a given UNI $u1$ agreed per [R5], the SP **MUST** monitor and report performance for the Ingress and Egress Application Flows identified by the 3-tuple $\langle afs, z, u1 \rangle$, for each of the following 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. This differs from the SD-WAN UNI to SD-WAN UNI definitions of performance attributes in MEF 70.1 [9]

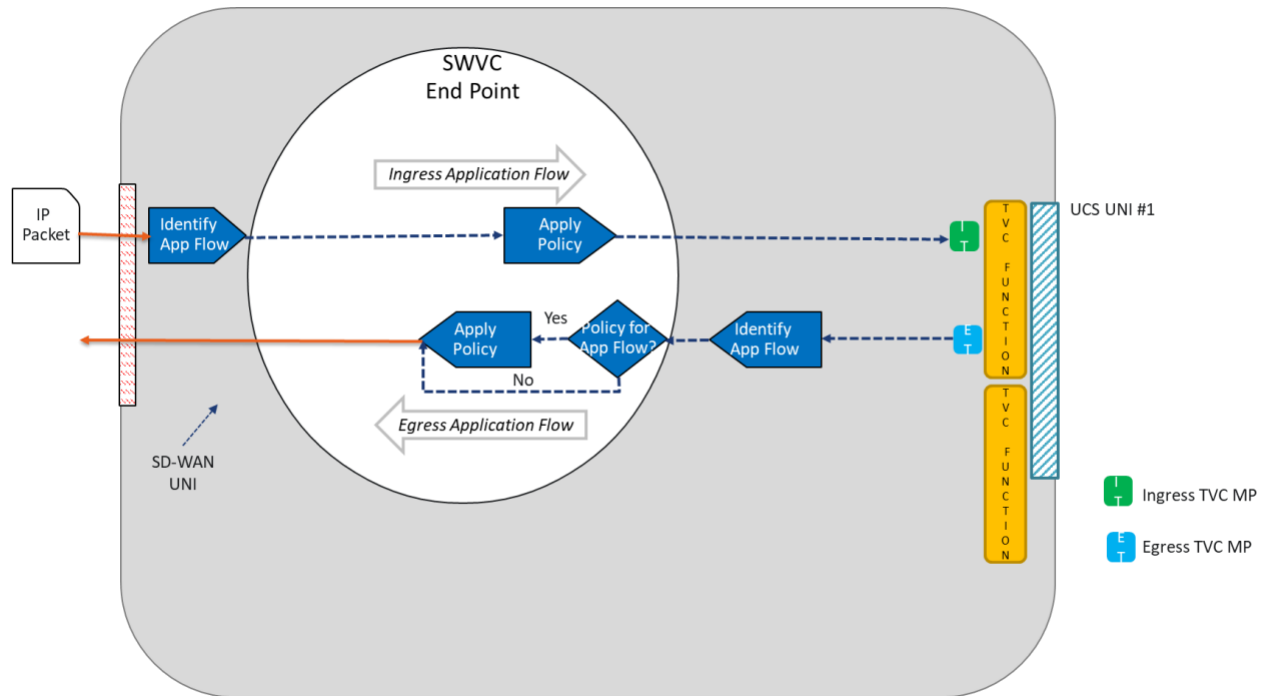


Figure 5 – TVC MP Locations

The TVC Functions in Figure 5 represent where the TVC originates/terminates. They are shown for informational purposes only.

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.

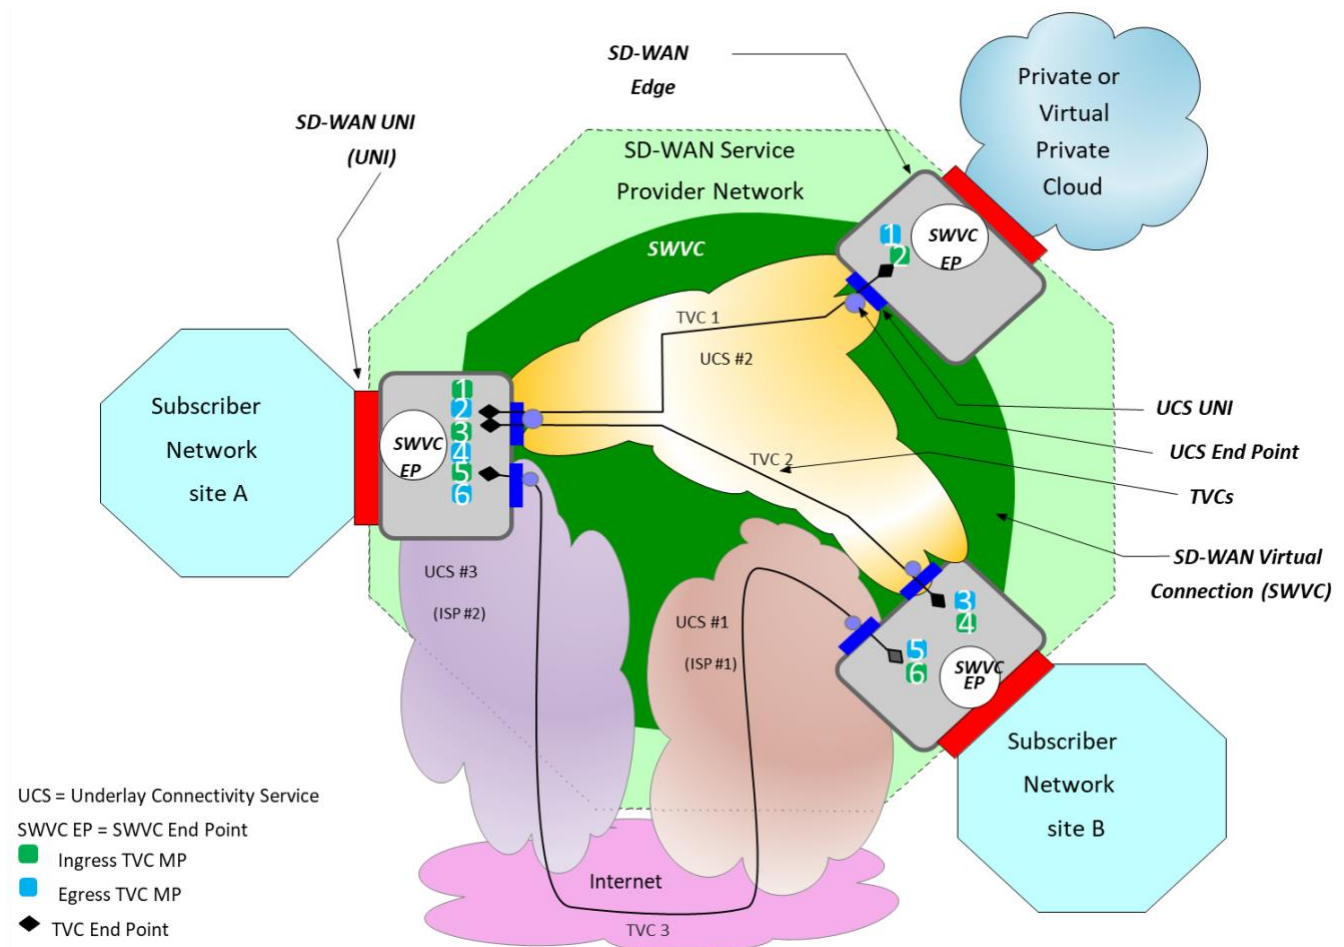


Figure 6 – Ingress TVC MP to Egress TVC MP ordered pair Example

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 PM Metric Calculation Interval Duration as specified in 7.2.2.2. This facilitates forwarding decisions based on performance.

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 Definitions

This section describes 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. 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 *<Application Flow Specification, Zone name, Ingress UNI Identifier, Egress UNI Identifier>* used for at least one of One-way Mean Packet Delay, One-way Mean 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 3-tuple *<Application Flow Specification, Zone name, UNI Identifier>* - used for at least one of Ingress MIR or 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 Mean 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.

Monitoring of Application Monitoring Flows and TVCs can be accomplished using synthetic IP Test Packets that pass over the Application Monitoring Flow or TVC that is being monitored. There are other mechanisms that might be used, such as hybrid monitoring where the Subscriber's IP Packets are modified to allow verification of performance.

Monitoring of Application Flows use the Subscriber IP Packets to perform byte counts which are then converted to Bits per second Measured Information Rate.

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 (Application Flow Specification, Zone name, Ingress UNI Identifier, Egress UNI Identifier). For Application Flows, it is the 3-tuple (Application Flow Specification, Zone name, UNI Identifier).

7.2.2.1 PM Metric List

For each Application Flow Specification, Zone name, the Subscriber and SP agree to a list of PM metrics that will be monitored. For each TVC, the Subscriber and SP agree if the three PM Metrics

will be monitored for the TVC. This list contains one or more PM Metrics as specified within this document.

7.2.2.1.1 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 received 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 $\delta_1 \dots \delta_n$.
 - Packets that straddle two intervals are discarded and not included in calculations.
- If Δ is *null* (no measurement values) then the One-way Mean Packet Delay for the Application Monitoring Flow over that interval is *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.
 - Packets that straddle two intervals are discarded and not included in calculations.
- Let Δ' be 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 section 9.8 of MEF 70.1 [9].
- If Δ' is *null* (no measurement values) then the value of the One-way Mean Inter-Packet Delay Variation for the PM Metric Calculation Interval is *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 value of 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} in Δ' .

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

- Let s represent the total number of IP Packets in the Application Monitoring Flow 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 transmitted during the 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 value of the One-way Packet Loss Ratio is 0.¹
 - If $s>0$, then the value of the One-way Packet Loss Ratio is $(s-r)/s$
 - Packets that straddle two intervals are discarded and not included in calculations.

[R9] The One-way Packet Loss Ratio **MUST** be represented as a percentage.

7.2.2.1.2 PM Metrics for Application Flows

Application Flow Ingress Measured Information Rate in bits per second 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 in bits per second for a given Egress Application Flow and a given PM Metric Calculation Interval is defined as:

- 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 [6] 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”.

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

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 the duration of the PM Metric Calculation Interval times the synthetic-rate qualified packets.

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

- [R10] 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 [9] and shown in Figure 3, is applied when calculating Application Flow Ingress Measured Information Rate.
- [R11] 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 [9] and shown in Figure 3, is applied when calculating Application Flow Egress Measured Information Rate.

7.2.2.1.3 PM Metrics for TVCs

The definitions for the PM Metrics for TVCs are detailed below. The impact of TVC PM Metric 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.
 - Packets that straddle two intervals are discarded and not included in calculations.

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. Then the One-way Mean Packet Delay for TVC 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 $\Delta = \{\kappa_1, \kappa_2, \kappa_3, \dots, \kappa_n\}$ 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.

- Packets that straddle two intervals are discarded and not included in calculations.
- Let Δ' be 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 section 9.8 of MEF 70.1 [9].
- If Δ' is *null*, then the value of 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 TVC One-way Mean Packet Delay Variation 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 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.
 - Packets that straddle two intervals are discarded and not included in calculations.
- 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 $(s-r)/s*100$
 - Packets that straddle two intervals are discarded and not included in calculations.

7.2.2.2 PM Metric Calculation Intervals

For each Monitored Entity, each PM Metric is calculated based on PM measurements taken over a sequence of contiguous time periods, all of equal duration – in other words, over a sequence of time intervals starting at $t + kP$ and each lasting for P , where t is a start time, P is a duration, and $k=\{0, 1, 2, \dots\}$. Each such time period is referred to as a PM Metric Calculation Interval. This is shown in Figure 7.

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 PM Metric Calculation Interval Duration*synthetic-rate qualified packets.




Figure 7 – PM Metric Calculation Interval Concept

For simplicity, all PM Metric Calculation Intervals for a given PM Metric have the same duration, when taking measurements for Application Flows or Application Monitoring Flows. When taking measurements for TVCs, all PM Metric Calculation Intervals have the same duration (regardless of PM Metric).

- [R12]** The sequences of PM Metric Calculation Intervals over which a given PM Metric is calculated for Application Flows or Applications Monitoring Flows **MUST** all use the same duration.
- [R13]** The sequences of PM Metric Calculation Intervals over which all PM Metrics are calculated for TVCs **MUST** all use the same duration.

[R12] and [R13] mean there are 6 possible PM Metric Calculation Interval Durations, which are agreed between the Subscriber and the SP.

- [R14]** The Subscriber and the SP **MUST** agree on the value of the following six PM Metric Calculation Interval Durations:
- PM Metric Calculation Interval Duration for One-way Mean Packet Delay for Application Monitoring Flows
 - PM Metric Calculation Interval Duration for One-way Mean Inter Packet Delay Variation for Application Monitoring Flows
 - PM Metric Calculation Interval Duration for One-way Packet Loss Ratio for Application Monitoring Flows
 - PM Metric Calculation Interval Duration for Ingress Measured Information Rate for Application Flows
 - PM Metric Calculation Interval Duration for Egress Measured Information Rate for Application Flows

- PM Metric Calculation Interval Duration for all PM Metrics for TVCs

These are shown in Figure 8.

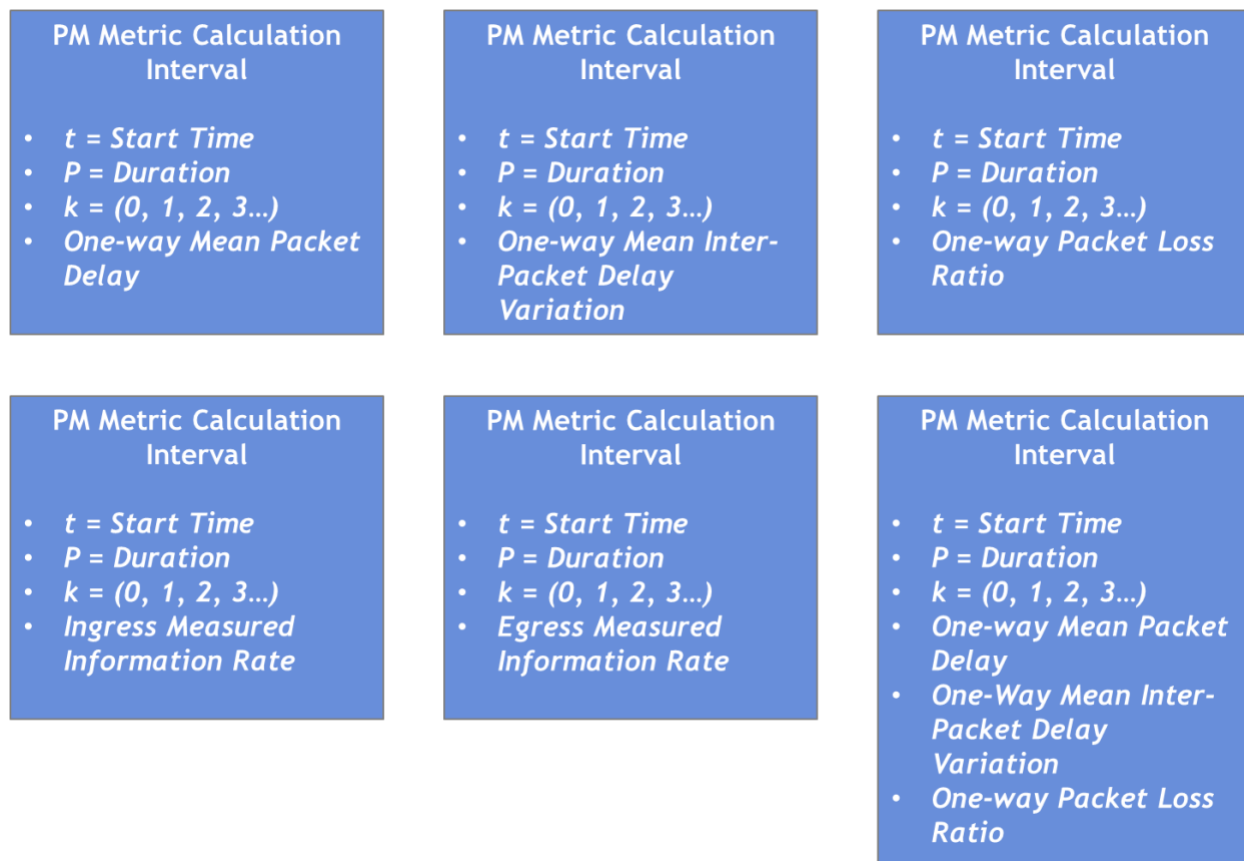


Figure 8 – PM Metric Calculation Durations

Again, for simplicity, all PM Metric Calculation Intervals with the same duration start at the same time. Thus, each of the six durations gives rise to a contiguous sequence of PM Metric Calculation Intervals, each of the specified duration. There are therefore six sequences of PM Metric Calculation Intervals, and every PM Metric calculation uses one of these six sequences. Note that different PM Metrics being measured on the same Monitored Entity (other than TVCs) might be calculated over different sequences of PM Metric Calculation Intervals.

- [R15] PM Metric values **MUST** be calculated based on measurements taken during a PM Metric Calculation Interval of the agreed duration.
- [R16] PM Metric Calculation Intervals **MUST** be contiguous.
- [R17] Sequences of PM Metric Calculation Intervals for different Monitored Entities that use a given duration per [R14] **MUST** start at the same time.
- [R18] A PM Metric Calculation Interval Duration **MUST** be less than or equal to 10000 milliseconds.

7.3 PM Metric Calculation and Reporting

This section describes the calculation and reporting of PM Metrics by the SD-WAN SP. Figure 9 illustrates the high-level sub-processes that are included in PM Metric Calculation and Reporting.

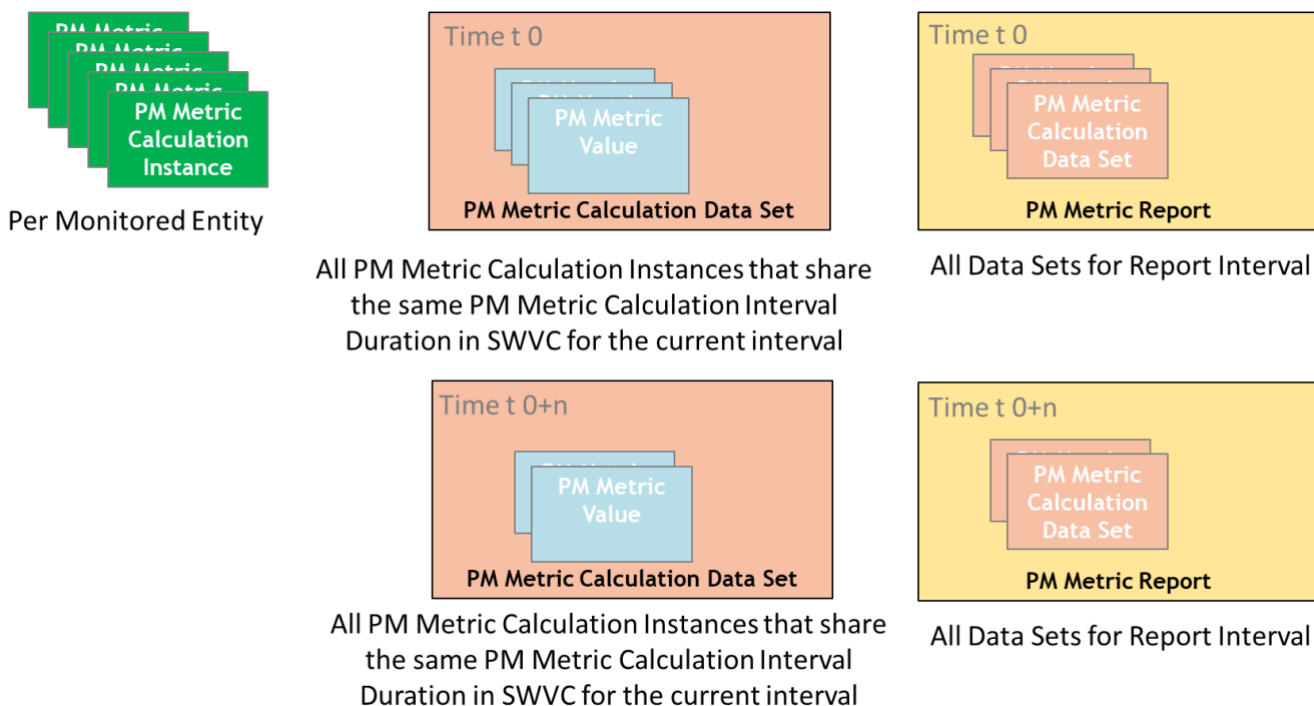


Figure 9 – PM Metric Calculation and Reporting Example

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 Interval Durations are aligned among different Monitored Entities i.e., they use the same PM Metric Calculation Interval Duration. The initial PM Metric Value for a given Monitored Entity may be calculated, but not reported, over a time period less than the specified PM Metric Calculation Interval 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. There is a sequence of PM Metric Calculation Data Sets one for each PM Metric Calculation Interval Duration. In the example above, there are two different PM Metric Calculation Interval Duration values being used which results in two sequences of 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 Duration
- List of pairs each consisting of PM Metric Calculation Instance 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 time and the duration of the PM Metric Report.

[R19] The Subscriber and SP **MUST** agree on the value of the PM Metric Report Interval.


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

The PM Metric Report Interval Duration is defined as the duration of the PM Metric Report 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.



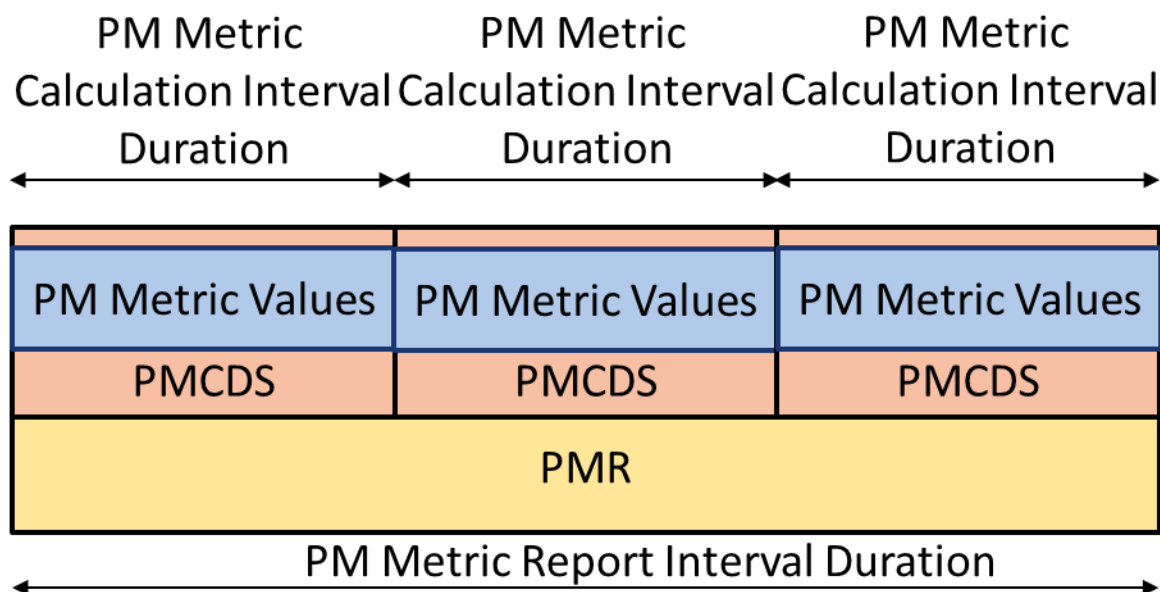
PMCDs– PM Metric Calculation Data Set

PMR – PM Metric Report

Figure 10 – PM Metric Calculation Interval Duration Equals PM Metric Report Interval Duration

Figure 10 shows an example where the value of the PM Metric Calculation Interval Duration and the value of the PM Metric Report Interval duration are equal, and the start times are aligned. All PM Metric Calculation Instances that use that duration are included in the report.

The value of the PM Metric Calculation Interval Duration and the value of the PM Metric Report Interval Duration do not have to be equal. Figure 11 shows an example of this.



PMCDS – PM Metric Calculation Data Set
PMR – PM Metric Report

Figure 11 – PM Metric Report Interval Greater Than PM Metric Calculation Interval Duration

In this example, there are three PM Metric Calculation Data Sets in each PM Metric Report. Multiple duration PM Calculation Data Sets can be included in the same PM Metric Report.

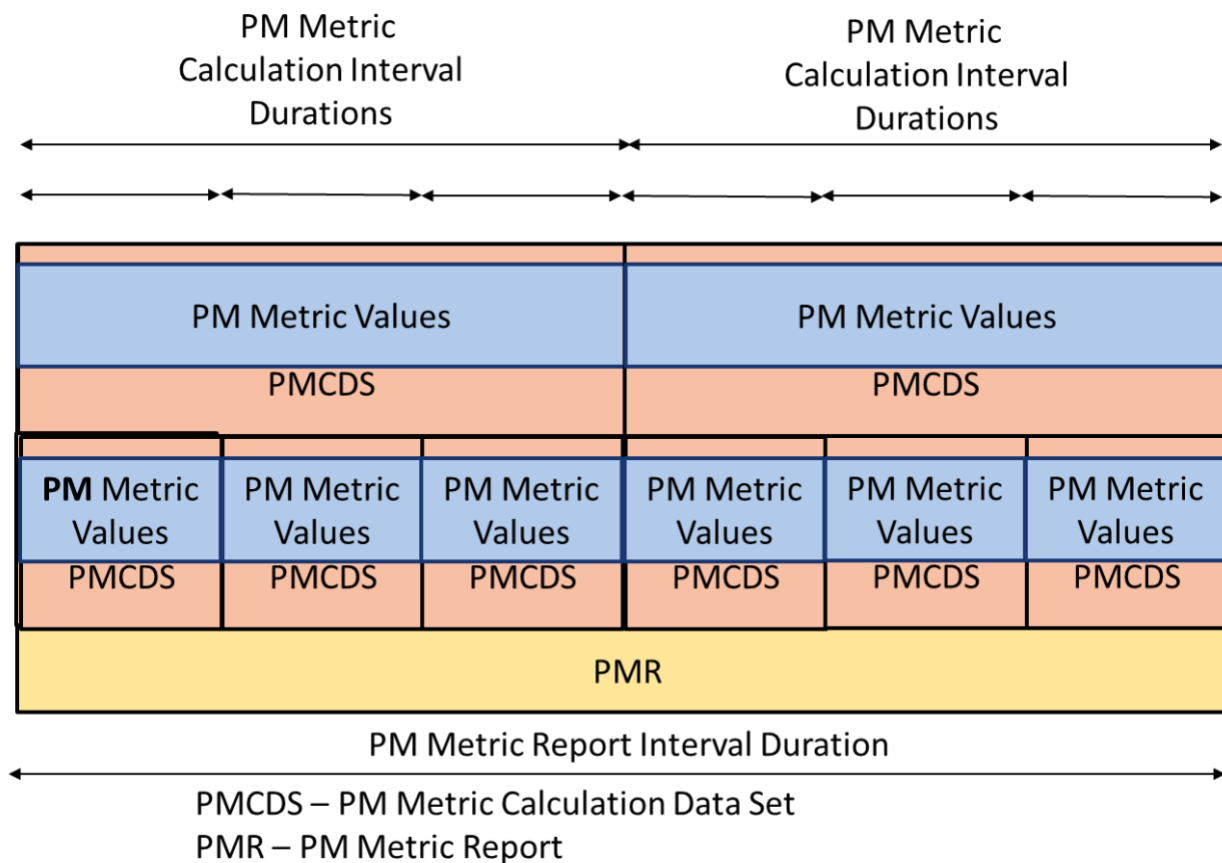


Figure 12 – Different PM Metric Calculation Interval Durations in a PM Metric Report

Figure 12 shows an example of when the value of the PM Metric Calculation Interval Duration is not equal for all PM Metrics. In this example, there are two different PM Metric Calculation Interval Durations that are included in the same PM Metric Report.

- [R20] The PM Metric Calculation Interval Durations **MUST** have their start times aligned to be included in the same PM Metric Report.
- [R21] The PM Metric Report Interval 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.

The retrieval of historical PM Metric Values, including after an PM Metric Calculation Instance is deleted, is left to the implementation.

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	Point in time in UTC	
PM Metric Calculation Data Set Duration	The length of time of the PM Metric Calculation Data Set Duration	Length of time	This is the PM Metric Calculation Interval Duration for the instances in the data set.
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.	One or more of with the value of that PM Metric: <ul style="list-style-type: none"> One-way Mean Packet Delay One-way Mean Inter-Packet Delay Variation One-way Packet Loss Ratio Ingress Measured Information Rate Egress Measured Information Rate 	This list contains values for PM Metric.

Table 8 – PM Metric Calculation Data Set Attributes

[R22] An SD-WAN PM Implementation **MUST** include the values of the attributes defined in Table 8 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 9.

Attribute	Value	Value Constraint	Comments
PM Metric Report Interval Start Time	The start time of the PM Metric Report interval.	Point in time in UTC	
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	One or more of with the value of that PM Metric: <ul style="list-style-type: none"> One-way Mean Packet Delay One-way Mean Inter-Packet Delay Variation One-way Packet Loss Ratio Ingress Measured Information Rate Egress Measured Information Rate 	Including all PM Metrics and PM Metric Values per PM Metric Calculation Instance

Table 9 – PM Metric Report Attributes

- [R24] The SD-WAN SP **MUST** provide a PM Metric Report that contains the attributes specified in Table 9.
- [R25] If an aggregated value of a PM Metric over the PM Metric Report Interval is reported (i.e., PM Metric calculated over multiple PM Metric Calculation Intervals), 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.

- [D1] 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 Duration 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 Mean 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 process of generating TCAs for a given PM Metric, a given PM Metric Calculation Instance, and a given set of TCA parameters. 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 7.4.3.3.

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.

Stateful and Stateless TCA Reporting are explained in the sections 7.4.1 and 7.4.2.

7.4.1 Stateful Threshold Crossing Alert Reporting

When using Stateful TCA reporting, each TCA Function has a SET TCA Window Threshold value and a CLEAR TCA Window Threshold value.

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 in Figure 13. The algorithm is defined in the requirements in section 7.4.3.2.

The TCA Function begins in the 'clear' state and stays in the 'clear' state for the first PM Metric Calculation Interval and all subsequent PM Metric Calculation Intervals until such time as the PM Metric Value has exceeded the Threshold. 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 value of 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-TCA's are generated while in this state.

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 the value of the 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.

Note that a SET-TCA can only be followed by a single CLEAR-TCA and a CLEAR-TCA can only be followed by a single SET-TCA.

The Stateful TCA algorithm for determining the state of the TCA function is shown in the flowchart in Figure 13. 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 $n = W, W + 1, \dots$

$St(n)$ = TCA Function value for interval n ('set' or 'clear')

W = 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 $j = n, n - 1, n - 2, \dots, n - W + 1$

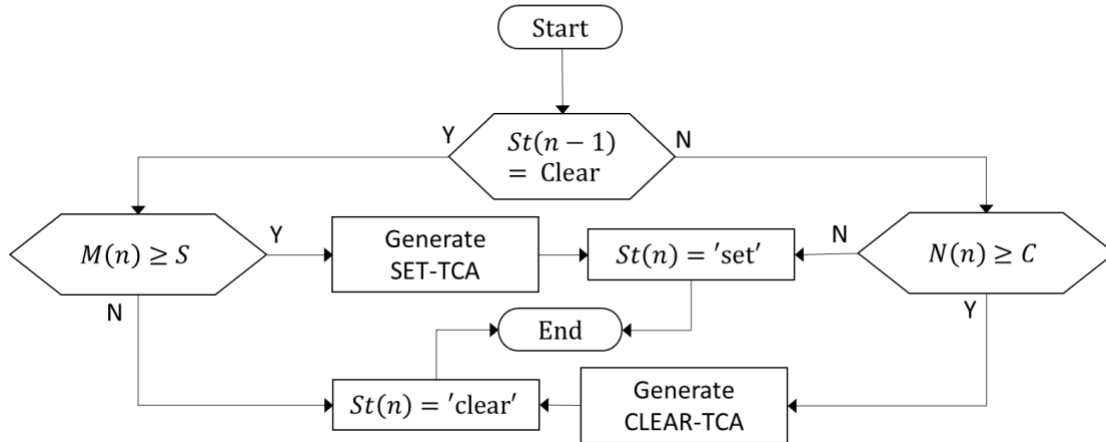


Figure 13 – Flowchart for Stateful TCA

The algorithm is formally defined in requirements [CR10]< and [CR11]< in section 7.4.3.2.

Figure 14 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 [CR9]< in section 7.4.3.2). 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 14 – Stateful 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 Threshold Value in PM Metric Calculation Intervals 13, 14, and 15 (i.e., 3 out of 6 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 Threshold Value in PM Metric Calculation Intervals 19, 20, 21, and 22 (i.e., 4 out of 6 which means that the 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.

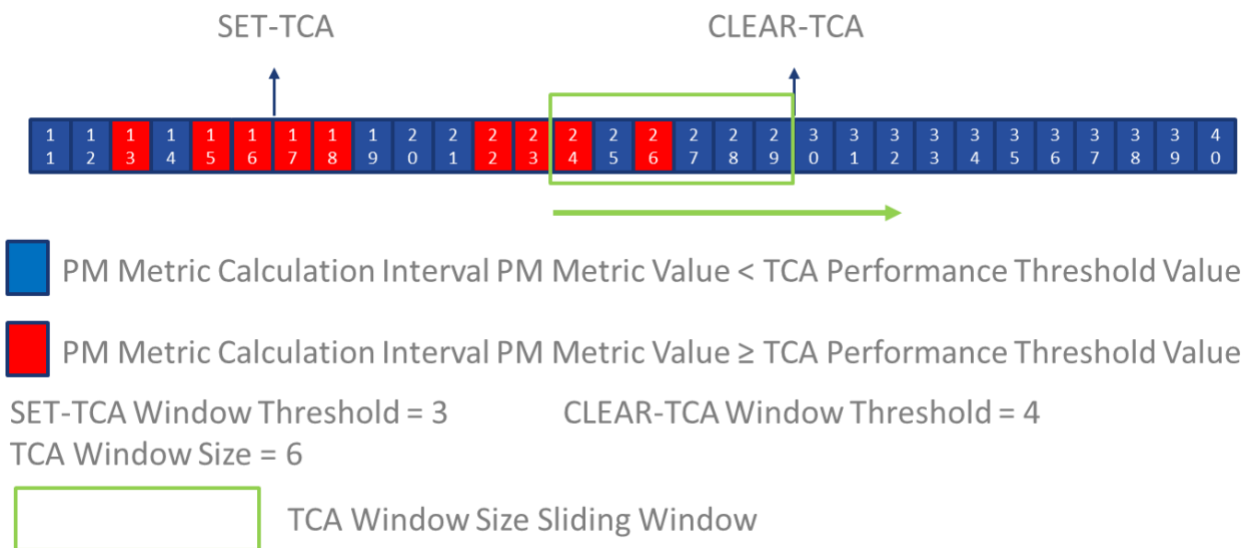


Figure 15 – Stateful TCA without Consecutive PM Metric Calculation Intervals ≥ TCA Performance Threshold Value Example

Figure 15 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 16. 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 (PM Metric Calculation Intervals 25, 27-29), even though the PM Metric Value in PM Metric Calculation Intervals 24 and 26 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 SET TCA Window Threshold and the CLEAR TCA Window Size Threshold are not used.

The algorithm is informally described in the following paragraph and the flowchart in Figure 15. The algorithm is defined in the requirements [CR13] below in section 7.4.3.3.

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 16.

$$\text{For intervals } n = 1, 2, \dots$$

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




Figure 16 – Flowchart for Stateless TCA w/o Damping Factor

The algorithm for Stateless TCA with Damping Factor is informally described in Figure 17. 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 requirements, [CR13]<, [CR14]<, [CR16]<, [CR17]<, [CR18]<, [CR19]<, and [CR20]< in section 7.4.3.3.

$$\text{For intervals } n = 1, 2, \dots$$

$$d = \text{Damping Factor} \geq 2, c \text{ is a counter}$$

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

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

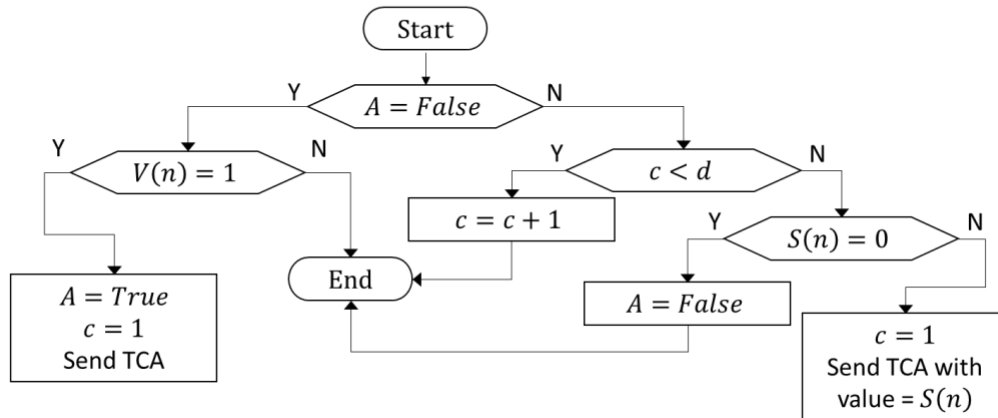


Figure 17 – Flowchart for Stateless TCA w/Damping Factor = *enabled*

Figure 18 shows an example of a Stateless TCA with the Damping Factor equal to *disabled*.




Figure 18 – Stateless TCA with Damping Factor = *disabled* Example

As shown in Figure 18, a TCA is generated for each PM Metric Calculation Interval in which a degraded condition exists. In this example, the TCA criterion is met in PM Metric Calculation Interval 13-22, 27, 33, 34, 36, and 38. TCAs can be generated in each of these PM Metric Calculation Intervals.

The Damping Factor is used to avoid generating a TCA for every PM Metric Calculation Interval with a degraded condition. The impact of the Damping Factor is shown in Figure 19.

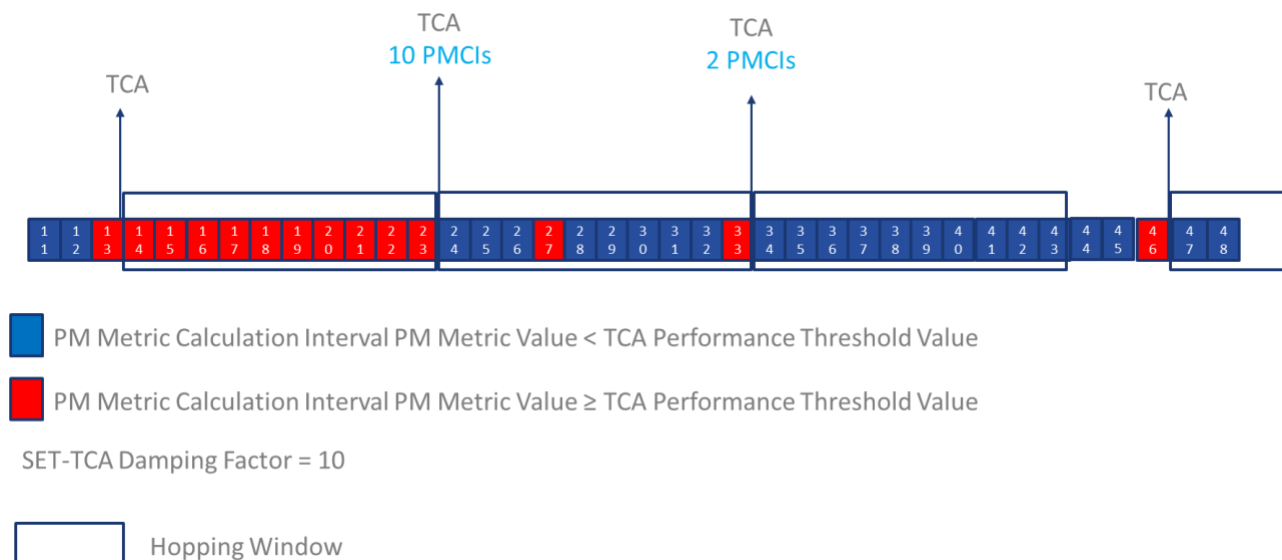


Figure 19 – Stateless TCA with Damping Factor Example

As shown in Figure 19, 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 value. 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 value 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. In this example, 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 19); this is a recommended parameter per [CD3]<. 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 19). 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 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 is deactivated. 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 Algorithm 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 46 which activates the Damping Factor again.

7.4.3 Threshold Crossing Alert Requirements

This section defines the requirements for the use of TCAs. There are three categories of requirements, Common that apply to both Stateful and Stateless, Stateful that only apply to Stateful TCAs, and Stateless that only apply to Stateless TCAs.

7.4.3.1 Common TCA Requirements

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

[D2] An SD-WAN PM Implementation **SHOULD** provide the ability to generate TCAs.

[CR1]<[D2] An implementation supporting TCAs **MUST** support at least one of Stateful or Stateless TCA Reporting.

[CR2]<[D2] An implementation supporting TCAs **MUST** support TCA Functions with different TCA Function parameter values for each Monitored Entity and each PM Metric.

[CD1]<[D2] An implementation supporting TCAs **SHOULD** support multiple TCA Functions with different TCA Function parameter values for a given Monitored Entity and PM Metric.

[CR3]<[D2] 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.

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.

[CR4]<[D2] An implementation supporting Stateful TCA Reporting **MUST** support SET-TCA functionality informally described in 7.4.1.

[CR5]<[D2] An implementation supporting Stateful TCA Reporting **MUST** support CLEAR-TCA functionality informally described in section 7.4.1.

[CR6]<[D2] An implementation supporting Stateful TCA Reporting **MUST** support the value of the SET-TCA Window Threshold being any integer value within a range of 1-300.

[CR7]<[D2] An implementation supporting Stateful TCA Reporting **MUST** support the value of the CLEAR-TCA Window Threshold being any integer value within a range of 1-300.

[CR8]<[D2] 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.

Note: 300 has been identified in this document as the maximum range size based on an assumption that PM Metric Calculations are occurring once per second. 300 is then equal to five minutes based on this assumption. It is believed that traffic rerouting may have taken place by this time.

Table 10 introduces symbols for the values of the TCA parameters. These symbols are used in requirements that follow the table.

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}
One-way Mean Inter-Packet	PT_{IPDV}	SWT_{IPDV}	CWT_{IPDV}	TWS_{IPDV}

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

Table 10 – Stateful TCA Reporting PM Metric Parameters

[CR9]<[D2] For an implementation supporting Stateful TCA Reporting, for each TCA Function, the sum of the values of the SET-TCA Window Threshold and the CLEAR-TCA Window Threshold **MUST** be greater than the value of the TCA Window Size.

[CR10]<[D2] An implementation supporting Stateful TCA Reporting **MUST** generate a SET-TCA notification message at the end of a PM Metric Calculation Interval and set the internal state of the TCA Function to ‘set’ if the TCA Function is in the ‘clear’ state and the calculated PM Metric value equals or exceeds the Performance Threshold in at least S of the last W PM Metric Calculation Intervals (including the one just ended), where S is the SET-TCA Window Threshold and W is the TCA Window Size..

[CR11]<[D2] An implementation supporting Stateful TCA Reporting **MUST** generate a CLEAR-TCA notification message at the end of a PM Metric Calculation Interval and set the internal state of the TCA Function to ‘clear’ if the TCA Function is in the ‘set’ state and the calculated PM Metric value is less than the Performance Threshold in at least C of the last W PM Metric Calculation Intervals (including the one just ended), where C is the CLEAR-TCA Window Threshold and W is the TCA Window Size.

[CR12]<[D2] 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.

Field Name	Field Value	Field Description
Monitored Entity	1-53 characters	Human readable text representing the Monitored Entity as described in section 3.
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 10.
Performance Monitoring Metric Value List	A list of the Numeric value of the PM Metric for the PM Metric Calculation Interval	PM Metric Values are reflected in ms for One-way Mean Packet Delay and One-way Mean Inter-Packet Delay Variation, and percentage for One-way Packet Loss Ratio
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 is used for loss.
Window Threshold Value	Numeric value as an Integer ≥ 1	The value of the SET-TCA or CLEAR-TCA Window Threshold.
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.
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. The criteria necessary for each value 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-CLEAR is shown as CLEARED.

7.4.3.3 Stateless TCA Reporting Requirements

The requirements in this section apply to Stateless TCA implementations.

- [CR13]<[D2] An implementation supporting Stateless TCA Reporting **MUST** support the TCA functionality informally defined in section 7.4.2.
- [CR14]<[D2] An implementation of Stateless TCA Reporting with the Damping Factor = *disabled* **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.
- [CD2]<[D2] An implementation of Stateless TCA Reporting **SHOULD** support the Damping Factor = *enabled*.
- [CR15]<[D2] An implementation of Stateless TCA Reporting that supports the Damping Factor **MUST** include a Damping Factor = *enabled/disabled* attribute.
- [CR16]<[D2] An implementation of Stateless TCA Reporting with the Damping Factor = *enabled* **MUST** support the Damping Factor integer values between 2-300 PM Metric Calculation Intervals.
- [CR17]<[D2] An implementation of Stateless TCA Reporting with the Damping Factor = *enabled* **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.
- [CR18]<[D2] An implementation of Stateless TCA Reporting with the Damping Factor = *enabled* **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.
- [CR19]<[D2] An implementation of Stateless TCA Reporting with the Damping Factor = *enabled* **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.
- [CR20]<[D2] An implementation of Stateless TCA Reporting with the Damping Factor = *enabled* **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.
- [CD3]<[D2] An implementation of Stateless TCA Reporting with the Damping Factor = *enabled*, 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.

[CR21]<[D2] An implementation of Stateless TCA Reporting **MUST** include the information shown in Table 12 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.
Monitored Entity	1-53 characters	Human readable text representing the Monitored Entity as described in section 3.
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 10.
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 is used for loss.
Performance Monitoring Metric Value	Numeric value as an Integer ≥ 1	The PM Metric Value for the PM Metric Calculation Interval duration in which the TCA Performance Threshold Value was equal to or exceeded by the PM Metric Value.
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. The criteria necessary for each value are agreed to by Subscriber and SP.

Table 12 – Stateless TCA Reporting Notification Message Fields

[CR22]<[D2] If an implementation of Stateless TCA Reporting has the Damping Factor = *enabled*, the Damping Factor information shown in Table 13 **MUST** be appended to the TCA notification message.

[CR23]<[D2] If an implementation of Stateless TCA Reporting includes reporting the number of PM Metric Calculation Intervals per [CD3]<[D2], 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 13 **MUST** be appended to the TCA Notification.

Field Name	Field Value	Field Description
Damping Factor	Numeric value	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	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 13 – Damping Factor TCA Notification Message Field

The Damping Factor can be used to reduce the number of TCAs generated by an implementation supporting 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 in Table 13 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 a Subscriber and/or the SD-WAN SP to perform additional research to determine 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 an SD-WAN Service to ensure it is ready for the Subscriber to begin using. It is performed to ensure that connectivity across UCSs exists between SD-WAN Edges. SRT for an 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 SD-WAN Service that is agreed by the Subscriber and SD-WAN SP to be Service Readiness Tested.

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 (the entity that is responsible for managing/orchestrating the SD-WAN Service), are configured with the basic SD-WAN Edge configuration used by the SD-WAN SP, and the 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 an SRT use case (section 8.1)
- A discussion of SRT terms and components (section 8.2)
- A description of Service Readiness Measurement Points (SRMPs), their location, and associated requirements (section 8.3)
- Tables that define which UCS attributes are tested, and which are reported (section 8.4)
- Service Readiness Testing Parameters and Results (section 8.5)
- 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 Use Case

This use case describes the verification of ordered pair of UCS End Points connectivity. 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 20 – SRT Order Pairs of UCS End Points Connectivity Verification Use Case Example

Figure 20 shows the use case for the verification of connectivity between SD-WAN Edges. SRMPs (as described in section 8.2) are used 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.

If one or more of the agreed on ordered pairs of UCS End Points represent backup UCSs, the SD-WAN SP may be required to remove from service one or more active UCSs before testing the backup UCSs.

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. The SP might arrange downtime with the Subscriber to avoid disrupting Subscriber traffic that is using the same UCS. The desirability of doing this is beyond the scope of this standard.

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 from which test measurements are performed. 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. Testing is performed using only 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 defined in this section. The following figures show the location of SRMPs in relationship to processing functions within the SD-WAN Edge. 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 21.

Editor Note 2: Should the above paragraph be deleted and R29 and R30 modified to include “at the UCS UNI without being processed by the functions associated with the SWVC End Point.”

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.

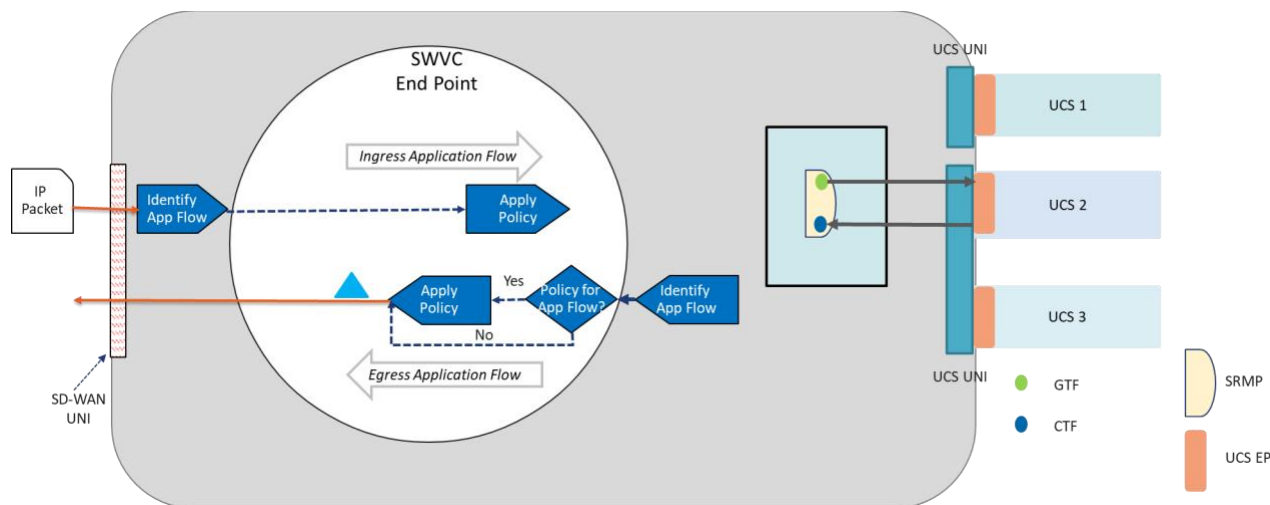


Figure 21 – SRMP Location of UCS End Points Connectivity Testing

[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 values of UCS Service Attributes (MEF 70.1 [9] section 12), UCS UNI Service Attributes (section 13), and UCS End Point Service Attributes (MEF 70.1 section 14) **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: As per [R33], the UCS Service Attributes are not required to be verified by SRT, however, the values of the attributes are required to be reported

8.5 Service Readiness Testing Parameters and Results

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

8.5.1 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 parameters for SRT are contained in a 4-tuple of the form $\langle OPairlist, TSC, Testn, Lostn \rangle$ where:

- *OPairlist* is a list of 2-tuples of the form $\langle EPid1, EPid2 \rangle$ where
 - *EPid1* is the value of the End Point identifier at the UCS UNI where the GTF is located
 - *EPid2* is the value of the End Point identifier at the UCS UNI where the CTF is located
 - *TSC* is the duration of the test for each of the entries in *OPairlist*
 - *Testn* is the number of IP packets generated uniformly distributed over *TSC* by the GTF for each of the entries in *OPairlist* during the test
 - *Lostn* is maximum number of lost IP packets for the test result to be *pass* for each entry in *OPairlist*
- [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 new ordered pairs of UCS End Points.
- [R36] When the SD-WAN SP is verifying connectivity, the Test Methodology defined in Table 14 in section 8.6.1 **MUST** be used.

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

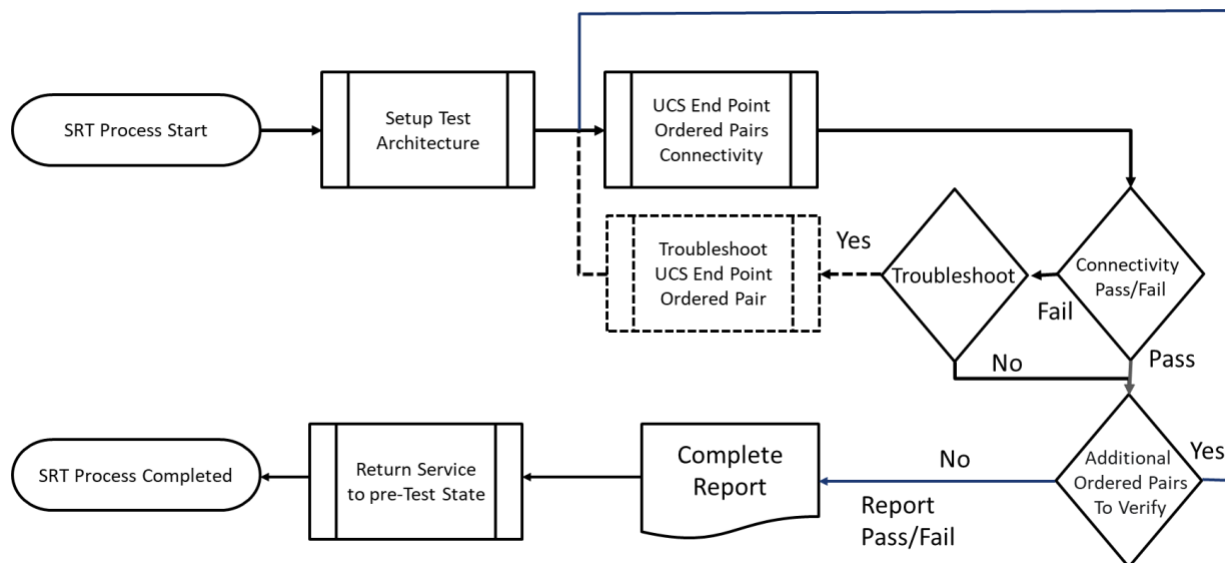


Figure 22 – UCS End Point Ordered Pairs Connectivity SRT Methodology

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 pairs as described in section 8.6.1. This step is repeated for each ordered pair that is agreed on to be included in SRT by the Subscriber and SD-WAN SP

The third step in the methodology is to report the UCS Service Attributes (see Complete Report step in Figure 22), 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 is shown by dashed lines. Whether troubleshooting takes place 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 14.

[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 21.. •SRMP₁ offers a number of IP Test Packets with the IP Destination Address 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 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 = the number of lost packets is less than or equal to the allowed value agreed to by Subscriber and SP Fail = the number of lost packets is greater than the allowed value
Remarks	None

Table 14 – UCS End Point Ordered Pair Connectivity Test Methodology

Editor Note 3: The Remarks row in the table is shown for completeness. At this time there are no Remarks, and the row will be deleted if none are added before LB.

[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 all ordered pairs between the SD-WAN Edges.

When the agreed 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] MEF 35.1, *Service OAM Performance Monitoring Implementation Agreement*, May 2015
- [5] MEF 48.1, *Carrier Ethernet Service Activation Testing*, November 2019
- [6] MEF 61.1, *IP Service Attributes*, May 2019
- [7] MEF 66, *Service OAM for IP Services*, July 2020
- [8] MEF 67, *Service Activation Testing for IP Services*, December 2020
- [9] 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 Attribute values. 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 for an SD-WAN Service. 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 Attribute values.