



MEF 105

**Performance Monitoring and Service
Readiness Testing for SD-WAN Service**

April 2024

Disclaimer

© MEF Forum 2024. 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 standard. MEF is not responsible or liable for any modifications to this standard made by any other party.

The receipt or any use of this standard 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 standard.

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	7
5	Conventions	8
5.1	Numerical Conventions	8
5.2	Diagram Conventions	8
6	Introduction	10
7	SD-WAN Performance Monitoring	13
7.1	SD-WAN Performance Monitoring Framework	13
7.2	Performance Metrics and Performance Monitoring	14
7.2.1	Application Flow Monitoring	15
7.2.2	TVC Monitoring	21
7.2.3	PM Definitions.....	24
7.2.3.1	<i>Monitored Entity</i>	24
7.2.3.2	<i>Performance Metric Calculation Intervals</i>	25
7.2.3.3	<i>Performance Metric Calculation Instance</i>	26
7.2.4	Performance Metric Definitions	27
7.2.4.1	<i>Performance Metrics for Application Monitoring Flows</i>	27
7.2.4.2	<i>Performance Metrics for Application Flows</i>	29
7.2.4.3	<i>Performance Metrics for TVCs</i>	29
7.3	PM Metric Calculation and Reporting.....	31
7.3.1	Performance Metric Calculation Data Set	37
7.3.2	Performance Metric Report.....	38
7.4	Threshold Crossing Alert.....	39
7.4.1	Stateful Threshold Crossing Alert Reporting	42
7.4.2	Stateless Threshold Crossing Alert Reporting.....	45
7.4.3	Threshold Crossing Alert Requirements.....	49
7.4.3.1	<i>Common TCA Requirements</i>	49
7.4.3.2	<i>Stateful TCA Requirements</i>	49
7.4.3.3	<i>Stateless TCA Reporting Requirements</i>	53
8	Service Readiness Testing for an SD-WAN Service	57
8.1	Verifying ordered pairs of UCS End Points Connectivity Use Case	57
8.2	Service Readiness Testing Terms and Components	59
8.3	Service Readiness Measurement Point Locations	59
8.4	UCS Service Attribute Reporting	60
8.5	Service Readiness Testing Parameters	60
8.6	Test Methodology	61
8.7	Test Record.....	65
9	References	66
Appendix A Service Readiness Testing vs Service Activation Testing (Informative)		67



Appendix B Acknowledgements (Informative) 68



List of Figures

Figure 1 – Diagram Conventions	9
Figure 2 – AFMP Locations for Two Application Monitoring Flows.....	15
Figure 3 – Measured Information Rate Byte Counter Locations.....	16
Figure 4 – Ingress AFMP to Egress AFMP ordered pair Example	17
Figure 5 – TVC MP Locations.....	22
Figure 6 –TVC MP Ordered Pair Example.....	23
Figure 7 – Performance Metric Calculation and Reporting Example.....	32
Figure 8 – Performance Metric Calculation Interval Duration Equals Performance Metric Report Interval Duration	35
Figure 9 – Performance Metric Report Interval Greater Than Performance Metric Calculation Interval Duration.....	36
Figure 10 – Different Performance Metric Calculation Interval Durations in a Performance Metric Report	37
Figure 11 – Flowchart for Stateful TCA.....	43
Figure 12 – Stateful TCA with Consecutive Performance Metric Calculation Intervals \geq TCA Performance Threshold Value Example	44
Figure 13 – Stateful TCA without Consecutive Performance Metric Calculation Intervals \geq TCA Performance Threshold Value Example	45
Figure 14 – Flowchart for Stateless TCA when Damping Factor Value = 1	46
Figure 15 – Stateless TCA with Damping Factor = 1 Example	46
Figure 16 – Flowchart for Stateless TCA w/Damping Factor \geq 2	47
Figure 17 – Stateless TCA with Damping Factor = 10 Example	48
Figure 18 – SRT Order Pairs of UCS End Points Connectivity Verification Use Case Example	58
Figure 19 – SRMP Location of UCS End Points Connectivity Testing	60
Figure 20 – UCS End Point Ordered Pairs SRT Process Flow.....	62



List of Tables

Table 1 – Terminology.....	5
Table 2 – Abbreviations.....	6
Table 3 – Numerical Conventions	8
Table 4 – Application Flow Specification, Zone name and Performance Metric Agreement Example	18
Table 5 – Monitored SD-WAN UNI ordered pair Agreement Example	18
Table 6 – Set of Monitored Measured Information Rate at SD-WAN UNI Agreement Example	19
Table 7 – Application Flow Specification, Zone name, Ordered Pair, and Performance Metric Example	20
Table 8 – Application Flow Specification Zone Name, UNI Identifier, and Performance Metric Overview Example.....	20
Table 9 – Performance Metric Calculation Instance Attributes.....	27
Table 10 – Performance Metric Calculation Data Set Attributes	38
Table 11 – Performance Metric Report Attributes	39
Table 12 – TCA Function Parameters	41
Table 13 – Stateful TCA Notification Message Fields	53
Table 14 – Stateless TCA Reporting Notification Message Fields	55
Table 15 – Damping Factor TCA Notification Message Field.....	56
Table 16 – UCS End Point Ordered Pair Connectivity Test Methodology	64



1 List of Contributing Members

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

- Bell Canada
- Ciena
- Cisco
- Orange



2 Abstract

This standard specifies definitions and requirements for Performance Monitoring and Service Readiness Testing for the SD-WAN Service described in MEF 70.1 [10]. Performance Monitoring consists of three parts:

- Definition and requirements for Performance Metrics and their associated measurement reference points that monitor Tunnel Virtual Connections and Application Flows,
- Definitions and requirements for calculating and reporting of the values for the Performance Metrics by the Service Provider to the Subscriber, and
- Definition and requirements for calculating and reporting by the Service Provider to the Subscriber of Threshold Crossing Alerts that capture performance degradations.

Service Readiness Testing consists of two parts:

- Definitions and requirements for testing the connectivity across Underlay Connectivity Services that exist between SD-WAN Edges and
- Requirements for the reporting of the Service Readiness Testing results and parameters

Note: This document is formatted to be read online and not to be printed. The text in some figures may not be optimized for a printed document.



3 Terminology and Abbreviations

This section defines the terms used in this standard. In many cases, the normative definitions to terms are found in other standards. 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 [6], MEF 61.1 [7], MEF 66 [8], and MEF 70.1 [10] are included in this standard by reference and are not repeated in the tables below.

Term	Definition	Reference
Application Flow Monitoring Point	A reference point that can be paired with another such reference point for the purpose of monitoring IP Packet performance of an Application Flow.	This document
Application Monitoring Flow	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 Performance Metric Calculation Intervals, within the TCA Window Size, for which the Performance 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 function for counting and discarding received IP Test Packets.	This document (Derived from MEF 67 [9])
Damping Factor	A parameter of Stateless TCA Reporting which is an integer > 0 that is used to suppress new TCAs.	This document
Egress Measured Information Rate	The Measured Information Rate at the Egress SD-WAN UNI.	This document
Egress Tunnel Virtual Connection Monitoring Point	The TVC MP at the egress side of the TVC.	This document
Generator Test Function	A function for generating and transmitting IP Test Packets.	This document (Derived from MEF 67 [9])
Ingress Measured Information Rate	The Measured Information Rate at the Ingress SD-WAN UNI.	This document



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Term	Definition	Reference
Ingress Tunnel Virtual Connection Monitoring Point	The TVC MP at the ingress side of the TVC.	This document
Measured Information Rate	The Information Rate expressed in bits per second measured during a single Performance Metric Calculation Interval for each Monitored Entity.	This document
Measurement IP Packet	Synthetic or Subscriber IP Packets that are used to measure performance.	This document
Monitored Entity	An Ingress or Egress Application Flow, an Application Monitoring Flow, or a TVC in a given direction.	This document
Performance Metric Calculation Data Set	A set of Performance Metric values for a given Performance Metric Calculation Interval.	This document
Performance Metric Calculation Instance	The measurement of a given Performance Metric, using the appropriate Performance Metric Calculation Interval Duration, for a given Monitored Entity.	This document
Performance Metric Calculation Interval Duration	The duration of the time intervals over which one or more Performance Metrics are calculated.	This document
Performance Metric Report	A report that contains the Performance Metric values for one or more Performance Metric Calculation Intervals.	This document
Performance Metric Report Interval Duration	The duration of each time interval over which one or more Performance Metric values from one or more Performance Metric Calculation Intervals are reported.	This document
SD-WAN Performance Monitoring	Performance Monitoring of the Application Flows, 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 standard for SD-WAN Performance Monitoring.	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 [6]
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



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Term	Definition	Reference
Service Readiness Testing Parameters	The parameters for SRT tests that are agreed between the SD-WAN SP and the Subscriber.	This document
SET-TCA Window Threshold	The number of Performance Metric Calculation Intervals, within the TCA Window Size, for which the Performance 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 [5]
Stateless TCA Reporting	A TCA reporting mechanism that generates a TCA when an alertable condition is detected.	This document derived from MEF 35.1 [5]
TCA Function	The process of generating TCAs for a given Performance Metric Calculation Instance, and a given set of TCA Function parameters.	This document
TCA Performance Threshold	The Performance Metric value that is compared against, for each Performance Metric Calculation Interval, when determining whether to generate a TCA.	This document
TCA Window Size	The sliding window of the number of consecutive Performance 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 Performance Metric, a particular set of TCA Function parameters and a particular Monitored Entity	This document derived from MEF 35.1 [5]
Tunnel Virtual Connection Measurement Point	A reference point that can be paired with another such reference point for the purpose of measuring TVC performance.	This document

Table 1 – Terminology



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Term	Definition	Reference
AFMP	Application Flow Monitoring Point	This document
CTF	Collector Test Function	MEF 67 [9]
GTF	Generator Test Function	MEF 67 [9]
MIR	Measured Information Rate	This document
PM	Performance Monitoring	This document derived from MEF 35.1 [5]
SD-WAN PM	SD-WAN Performance Monitoring	This document
SD-WAN SRT	SD-WAN Service Readiness Testing	This document
SRMP	Service Readiness Measurement Point	This document
SRT	Service Readiness Testing	This document
TCA	Threshold Crossing Alert	This document derived from MEF 35.1 [5]
TVC MP	Tunnel Virtual Connection Measurement Point	This document
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 standard 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 standard uses the prefix notation to indicate multiplier values as shown in Table 3.

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 3 – Numerical Conventions

5.2 Diagram Conventions

The diagrams in this standard have several components that appear frequently. These components are shown in Figure 1:

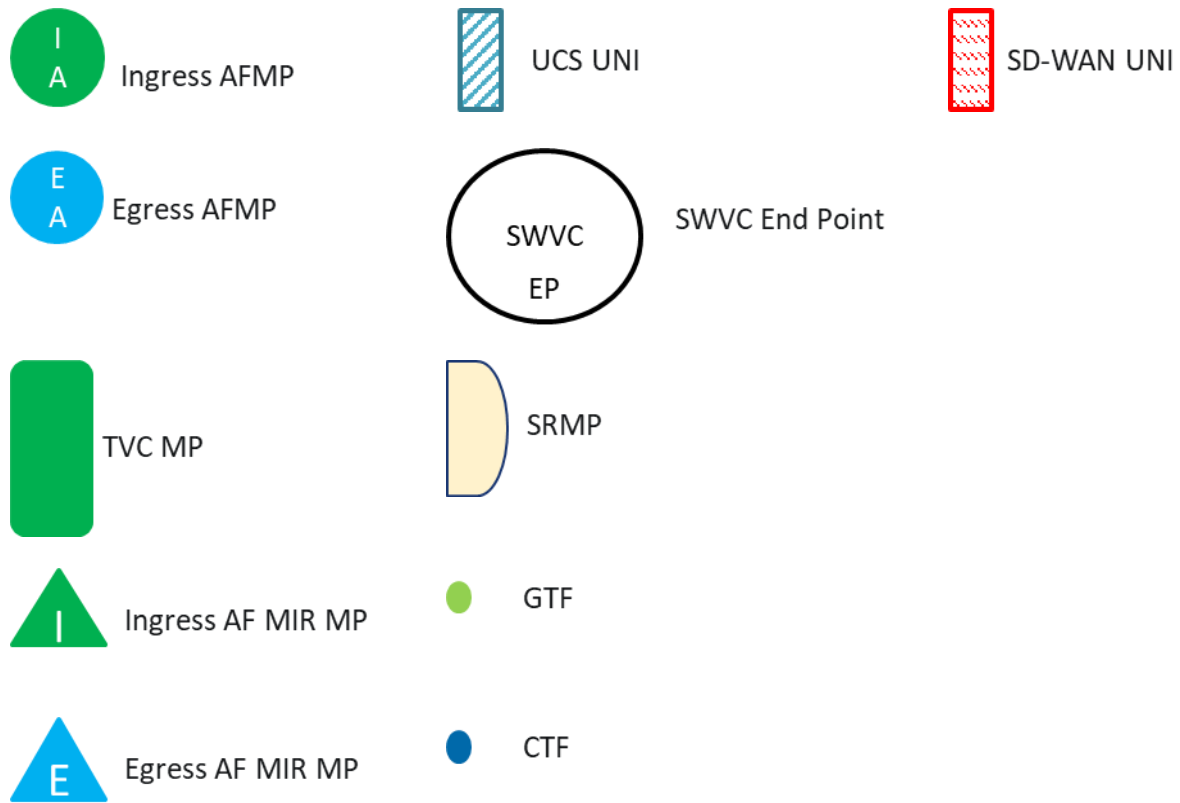


Figure 1 – Diagram Conventions



6 Introduction

This standard describes the requirements for PM in section 7 and SRT in section 8 pertinent to an SD-WAN Service as described and defined in MEF 70.1 [10].

For the SD-WAN Service standardized in MEF 70.1 [10], 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 document 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) described in this standard provides a documented method for monitoring the performance of the service. SD-WAN Service Readiness Testing (SD-WAN SRT) is defined as 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 standard 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: IP Packet is defined in MEF 61.1 [7] 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”.

Note that IP Packets in a given Ingress Application Flow might be directed towards multiple Egress UNIs. An Application Flow may therefore contain multiple Application Monitoring Flows each directed towards a different Egress UNI.

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

Measurements of Tunnel Virtual Connections (TVCs) are performed between pairs of TVC Measurement Points which are located between SWVC End Points and UCS UNIs within an SD-WAN Edge as described in section 7.

As described in section 7 of this standard, SD-WAN PM addresses the SD-WAN PM statistics and information that the SD-WAN SP can share with the SD-WAN Subscriber. These are:

- Measured Information Rate (Ingress and/or Egress) per Application Flow
- Performance Metric values per Application Monitoring Flow
- Performance Metric values per ordered pair of Tunnel Virtual Connection (TVC) end points
- Threshold Crossing Alerts (TCA) for the Performance Metric values.

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



Performance Monitoring and Service Readiness Testing for SD-WAN Service

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.

Performance 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 over short time periods (see section 7.1 for the specified time period). As an example of 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 SP can use these to determine if packets are passing over an Application Monitoring Flow and share this information with the Subscriber e.g., via a Subscriber Portal.

Some Performance Metrics can be used to make forwarding decisions within the SD-WAN Service. Performance 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 [10] 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 actual implementation of using these Performance Monitoring Metrics for forwarding purposes is beyond the scope of this document.

The SD-WAN SRT is described in Section 8 of this standard 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 [10].
- 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: the term Service Readiness Testing (SRT) is limited to the functions and metrics defined within this standard.

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 between UCS End Points is known to have existed at the time of SRT.



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Service Readiness Testing verifies whether 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 [10]. 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 standard 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 standard 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 SP unless otherwise noted.

7 SD-WAN Performance Monitoring

SD-WAN Performance Monitoring (SD-WAN PM) is defined as Performance Monitoring (PM) of the Application Flows, Application Monitoring Flows and TVCs in an SD-WAN Service. An SD-WAN SP may measure and report a variety of Performance Metrics describing the performance of the various components of the SD-WAN Service, this standard is limited to what is defined below. The tools used to measure SD-WAN performance are not defined in this standard.

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

- SD-WAN Performance Monitoring Framework defined in section 7.1
- Performance Metrics for Application Flows, Application Monitoring Flows, and TVCs defined in section 7.2
- Performance Metric Calculation and Reporting defined in section 7.3
- Threshold Crossing Alerts (TCAs) defined in section 7.4

The set of Application Flows and Application Monitoring Flows that are monitored are agreed-on by the SD-WAN SP and the Subscriber. This set can be a strict subset of all Application Flows and 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 Performance Metric values are the results of PM measurements and calculations performed using IP Packets. Many requirements in this standard apply to the SD-WAN PM Implementation used to perform SD-WAN PM.

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

7.1 SD-WAN Performance Monitoring Framework

SD-WAN PM starts with measurement intervals that are significantly shorter (10 seconds or less) as defined in [D1] 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. Performance Metric calculation results are used not only to measure the performance of Application Flows, Application Monitoring Flows, and TVCs but can also be used to make IP Packet forwarding decisions when PERFORMANCE criteria as described



in MEF 70.1 [10] are included in Policies. There are differences between how the Performance Metrics used for agreement of the PERFORMANCE policy criterion between the SP and Subscriber are defined in MEF 70.1, and the Performance Metrics used for measurement are defined in this standard. Resolution of these differences is beyond the scope of this standard.

The results of each SD-WAN measurement interval are reported based on the definition of the Performance Metric Report defined as a report that contains the Performance Metric values for one or more Performance Metric Calculation Intervals. The interval in which Performance Metric Reports are generated and the Performance Metrics included in the report are specified in the Performance Metric Report definition. Performance Metric Reports can be used by the SD-WAN SP, and they may be shared with the Subscriber. If shared with the Subscriber, the Subscriber and SP agree to what is shared.

In addition to Performance 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 standard).

7.2 Performance Metrics and Performance Monitoring

This section describes the measurement methods and Performance Metrics that are defined for SD-WAN Service. The Performance 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 can apply to Application Monitoring Flows as agreed by the Subscriber and the SP and are used for monitoring TVCs. Measured Information Rate can apply to Ingress and Egress Application Flows located at SD-WAN Edges as agreed by the Subscriber and SP.

Performance Metrics (One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, and One-way Packet Loss Ratio) defined in this standard are calculated from measurements taken using measurement points that are different than the measurement points described in MEF 70.1 [10] section 15. (see section 7.2.4 for details).

Application Monitoring Flow and TVC Performance Monitoring are described in section 7.2.1 and section 7.2.2 respectively. PM definitions are defined in section 7.2.3.

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 over either an Application Monitoring Flow or a TVC, are described in section 7.2.4. Measured Information Rate is also described in section 7.2.4.

7.2.1 Application Flow Monitoring

This section describes the monitoring of Application Flows and Application Monitoring Flows. The definitions of Performance Metrics values are contained in sections 7.2.4.1 and 7.2.4.2.

An Application Monitoring Flow is a set of IP Packets that match a given Ingress Application Flow and are directed towards a given Egress UNI. An Application Flow Monitoring Point is a reference point that can be paired with another such reference point for the purpose of monitoring IP Packet performance of an Application Flow. The location of Ingress and Egress Application Flow Monitoring Points is shown in Figure 2.

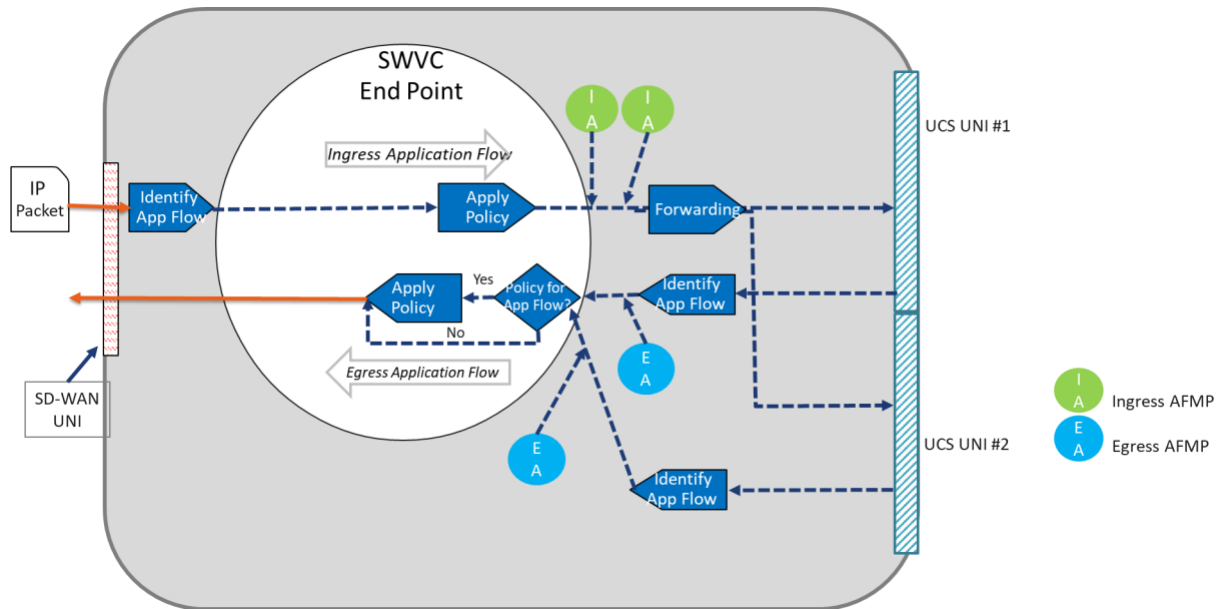


Figure 2 – AFMP Locations for Two Application Monitoring Flows

IP Packets in a given Ingress Application Flow might be directed towards multiple Egress UNIs. The location of the Application Flow Monitoring Points and Application Flow Byte Counters are discussed below:

- An Ingress Application Flow Monitoring Point (AFMP) is located after where any Policies are applied and before the Forwarding function (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 Performance Metrics, measurements are performed between the ordered pairs of AFMPs, <Ingress AFMP, Egress AFMP>, with the Ingress AFMP adjacent to an SWVC End Point and the Egress AFMP adjacent to another SWVC End Point.

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

Note: A Measured Information Rate of 0 Mb/s is considered a valid Performance Metric value. Implementations that do not use synthetic IP Packets or modify Subscriber IP Packets to perform measurements result in Performance Metric values that are reported as Undefined (see sections 7.2.4.1 and 7.2.4.3) if no subscriber IP Packets are present during the Performance Metric Calculation Interval.

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. The actual format of the Measurement IP Packets is beyond the scope of this document.

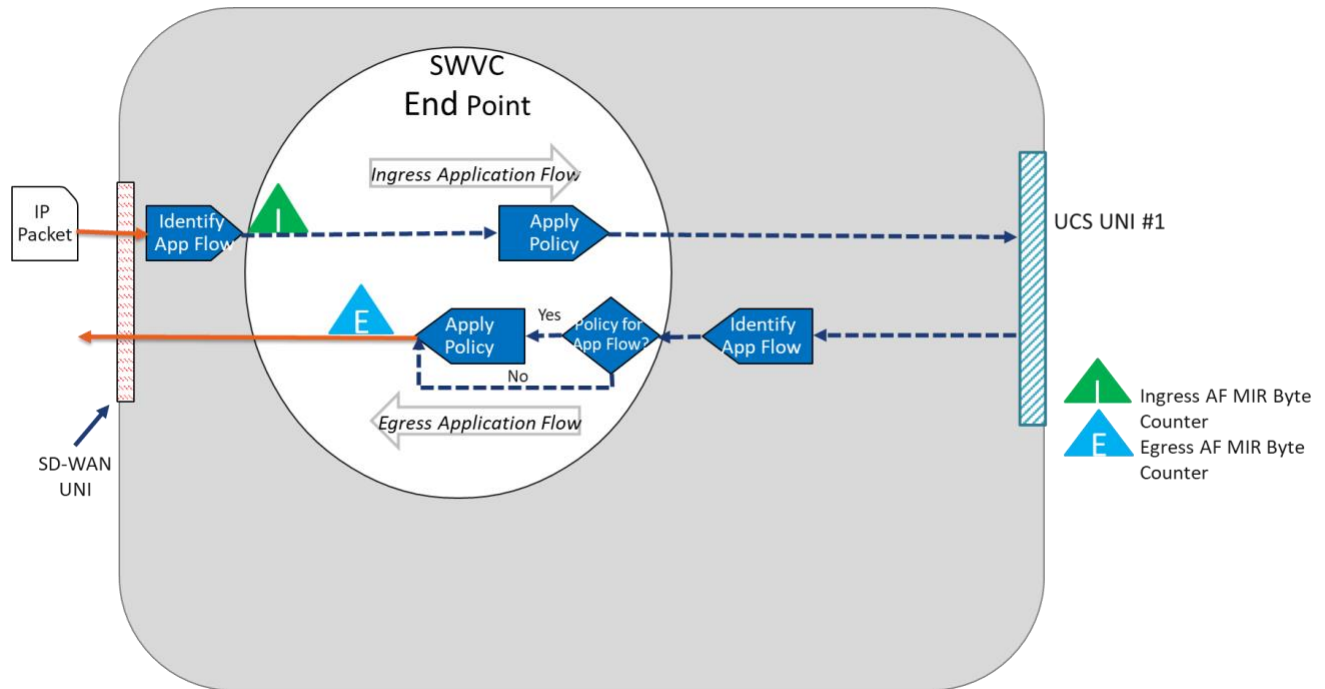


Figure 3 – Measured Information Rate Byte Counter Locations

Ingress Measured Information Rate is defined as the Measured Information Rate at the Ingress SD-WAN UNI.

Egress Measured Information Rate is defined as the Measured Information Rate at the Egress SD-WAN UNI.




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 Performance Metric values are calculated from the measurements. In the example, AFMP ordered pair 1 monitors Application Monitoring Flow 1, AFMP ordered pair 2 monitors Application Monitoring Flow 2, AFMP ordered pair 3 monitors Application Monitoring Flow 3, and AFMP ordered pair 4 monitor Application Monitoring Flow 4.

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 on the set of Performance Metrics to be monitored and reported by the SP to the Subscriber and the method by which they are reported.

The method used to report the metric values is outside the scope of this standard.

An example of the agreed Performance Metrics to be monitored is shown in Table 4.

Application Flow Specification, Zone name	Performance 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 4 – Application Flow Specification, Zone name and Performance Metric Agreement Example

Note: Performance Metrics are agreed to for each Application Flow Specification, Zone name combination. Different Performance Metrics can be agreed to for different Application Flow Specification, Zone name combinations.

[R3] The Subscriber and SP **MUST** agree on a single set of SD-WAN UNI ordered pairs which will be monitored as described in [R5].

An example of monitored SD-WAN UNI ordered pair agreement is shown in Table 5.

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

Table 5 – Monitored SD-WAN UNI ordered pair Agreement Example

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



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Note that the combination of Application Flow Specification and Zone name (from [R1]) and an ordered pair of UNIs (from [R3]) uniquely identifies an Application Monitoring Flow.

[R4] The Subscriber and SP MUST agree on a single set of SD-WAN UNIs at which Measured Information Rate will be monitored per [R6].

An example of a list of SD-WAN UNI identifiers is shown in Table 6.

SD-WAN UNI Identifiers
A
B
C
D

Table 6 – Set of Monitored Measured Information Rate at SD-WAN UNI Agreement Example

Note that the combination of Application Flow Specification and Zone Name (from [R1]) and a UNI (from [R4]) uniquely identifies an Ingress Application Flow and an Egress Application Flow. The SD-WAN UNI Identifiers for the ordered pairs agreed to for Application Flow Monitoring and the SD-WAN UNI Identifiers agreed to for Measured Information Rate monitoring do not have to be the same.

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 7 shows an example of the Performance 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 agreed to by the Subscriber and SP).

Application Flow Specification, Zone name	SD-WAN UNI ordered pair	Performance Metric
Voice, Corporate	A-B	One-way Mean Packet Delay One-way Inter-Packet Delay Variation One-way Packet Loss Ratio
Voice, Corporate	B-A	One-way Mean Packet Delay One-way Inter-Packet Delay Variation One-way Packet Loss Ratio
Voice, Corporate	A-D	One-way Mean Packet Delay One-way Inter-Packet Delay Variation One-way Packet Loss Ratio
Video Streaming, Corporate	A-B	One-way Mean Packet Delay



Performance Monitoring and Service Readiness Testing for SD-WAN Service

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

Table 7 – Application Flow Specification, Zone name, Ordered Pair, and Performance Metric Example

Given the examples above, Table 8 shows an example of the Performance 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	Performance Metric
Voice, Corporate	A	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	B	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	C	Ingress Measured Information Rate Egress Measured Information Rate
Voice, Corporate	D	Ingress Measured Information Rate Egress Measured Information Rate
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
Video Streaming, Corporate	A	Ingress Measured Information Rate 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

Table 8 – Application Flow Specification Zone Name, UNI Identifier, and Performance Metric Overview Example

[R5] For each pair of Application Flow Specification and Zone name, $\langle afs, z \rangle$, agreed per [R1], and each ordered pair of UNIs $\langle u1, u2 \rangle$ agreed per [R3], the SP **MUST** monitor performance for the Application Monitoring Flow $\langle afs, z, u1, u2 \rangle$, for each of the Performance Metrics that were agreed for the pair $\langle afs, z \rangle$ per [R1] from the following list of Performance Metrics:

- One-way Mean Packet Delay

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

[R6] For each pair of Application Flow Specification and Zone name, $\langle afs, z \rangle$, agreed per [R1], and each UNI $u1$ agreed per [R4], the SP **MUST** monitor the Application Flow Ingress Measured Information Rate and/or Application Flow Egress Measured Information Rate (as agreed for the pair $\langle afs, z \rangle$ per [R1]) for the Ingress and/or Egress Application Flows (respectively) identified by the 3-tuple $\langle afs, z, u1 \rangle$.

7.2.2 TVC Monitoring

This section describes the TVC monitoring requirements and the locations of the TVC Monitoring Points. The definitions of the Performance Metrics for TVCs are contained in section 7.2.4.3.

A TVC is defined as a forwarding relationship between two SD-WAN Edges. PM is performed unidirectionally on a TVC using an ordered pair of Tunnel Virtual Connection Measurement Points (TVC MP). A TVC MP is defined as a reference point that can be paired with another such reference point for the purpose of measuring TVC performance.

A TVC MP is located between the UCS UNI and the location inside the SD-WAN Edge where ingress and egress Policies are enforced. For One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, and One-way Packet Loss Ratio Performance Metrics, measurements are performed between the ordered pair of Ingress TVC MP located on one SD-WAN Edge and Egress TVC MP located on another SD-WAN Edge as shown in Figure 5.

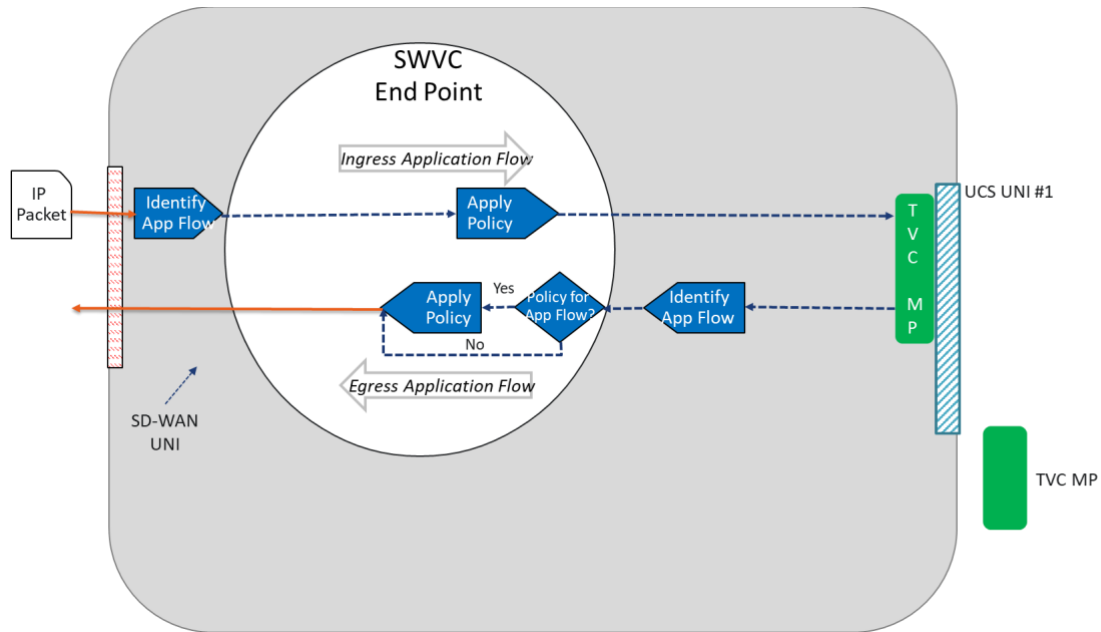


Figure 5 – TVC MP Locations

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

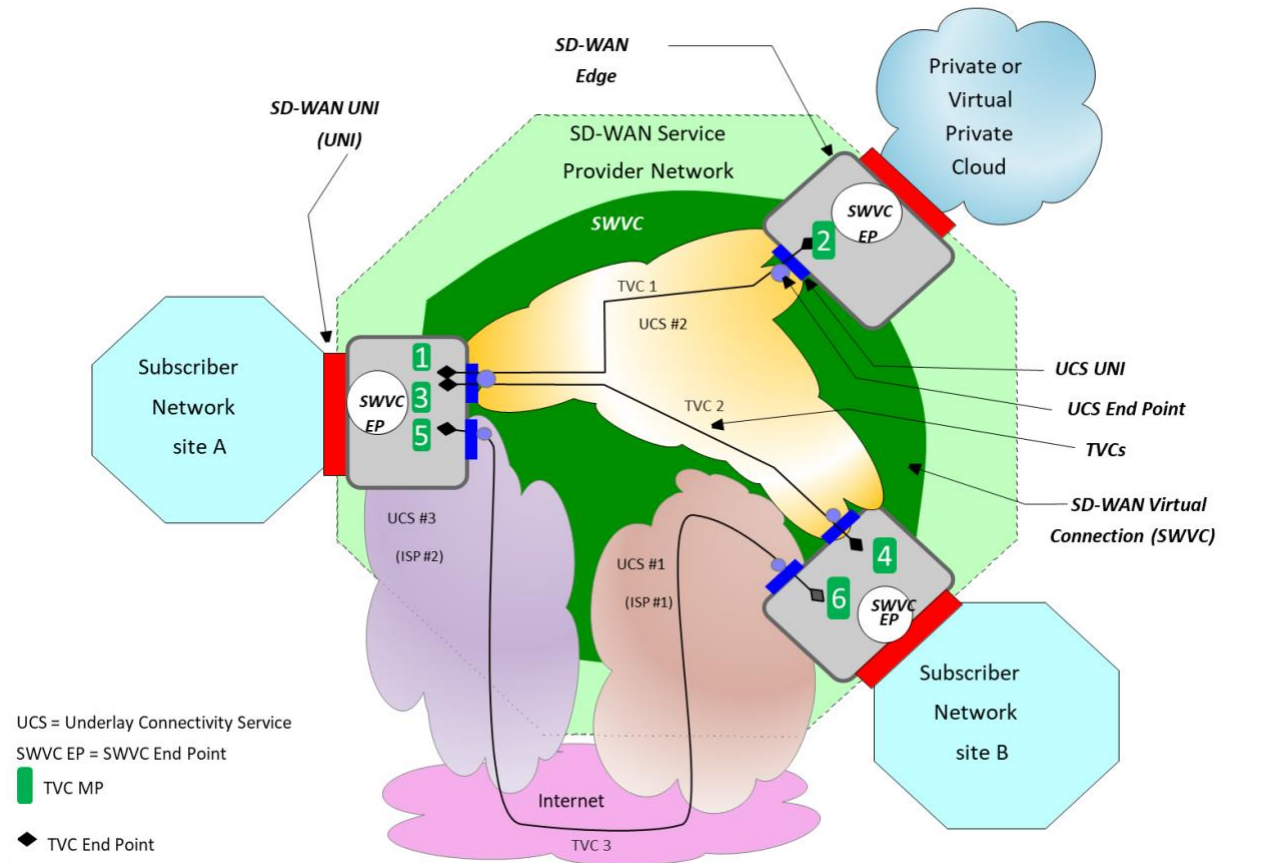


Figure 6 –TVC MP Ordered Pair Example

Figure 6 shows six TVC MP ordered pairs, identified as $\langle 1, 2 \rangle$, $\langle 2, 1 \rangle$, $\langle 3, 4 \rangle$, $\langle 4, 3 \rangle$, $\langle 5, 6 \rangle$, and $\langle 6, 5 \rangle$, that are monitoring three TVCs. Each ordered pair is made up of an Ingress TVC MP and an Egress TVC MP. IP Packets are generated or observed at the Ingress TVC MP and collected or observed at the Egress TVC MP. Performance Metric measurements are performed using these packets, and Performance 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 $\langle 1, 2 \rangle$ and $\langle 2, 1 \rangle$ monitor TVC 1, TVC MP ordered pairs $\langle 3, 4 \rangle$ and $\langle 4, 3 \rangle$ monitor TVC 2, and TVC MP ordered pairs $\langle 5, 6 \rangle$ and $\langle 6, 5 \rangle$ monitor TVC 3.

[R7] The SP **MUST** monitor all TVCs in an SWVC in both directions for the following Performance 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 Performance Metric Calculation Interval Duration, defined as the duration of the time intervals over which one or more Performance Metrics are calculated. This facilitates forwarding decisions based on performance.

Note: Monitoring of TVCs can provide the Subscriber with information about the performance of the overlay. In addition, Application Flow forwarding decisions based on PERFORMANCE criteria can be made by the Service Provider.

7.2.3 PM Definitions

This section describes the Monitored Entity, the Performance Metric Calculation Interval, and the Performance Metric Calculation Instance. For details on how one-way measurements can be performed please see MEF 66 [8] section 8.3.

7.2.3.1 Monitored Entity

A Monitored Entity is one of the following:

1. An Application Monitoring Flow flowing between a pair of UNIs, identified by the 4-tuple *<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 or One-way Packet Loss Ratio
2. An Application Flow 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
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: *UCS CoS Name* setting depends on the underlying UCS technology. 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*. When the UCS is Layer 3, the CoS Name is as described in MEF 61.1 [7]. When the UCS is Layer 2, the *UCS CoS Name* is as described in MEF 10.4 [4].

The set of Monitored Entities for an SWVC is based on the information agreed with the Subscriber per [R1], [R3] and [R4].

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 as defined in MEF 66 [8] where the Subscriber's IP Packets are modified to allow monitoring of performance.

MIR is calculated by measuring the byte counts of Application Flows. These byte counts are then converted to bits per second Measured Information Rate.

7.2.3.2 Performance Metric Calculation Intervals

For each Monitored Entity, each Performance Metric is calculated based on performance measurements taken over a sequence of contiguous time intervals 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 an arbitrary UTC time chosen by the Service Provider that is before the SD-WAN service is first provided to the Subscriber, P is a duration, and $k = 0, 1, \dots$. Each interval starting at $t + kP$ and ending at $t + (k + 1)P$ is referred to as a Performance Metric Calculation Interval.

The value of t applies to all Monitored Entities and Performance Metrics while the value of P can be different for different Monitored Entities and Performance Metrics as described below.

[R8] The sequences of Performance Metric Calculation Intervals over which a given Performance Metric is calculated for Application Flows or Applications Monitoring Flows **MUST** all use the same duration.

[R9] The sequences of Performance Metric Calculation Intervals over which all Performance Metrics are calculated for TVCs **MUST** all use the same duration.

The Performance Metric Calculation Intervals are therefore determined by the values of 6 durations as follows:

- P_D : The duration specified as a positive integer number of milliseconds used for calculating the One-way Mean Packet Delay Performance Metric applied to all Application Monitoring Flows,
- P_V : The duration specified as a positive integer number of milliseconds used for calculating the One-way Mean Inter-Packet Delay Variation Performance Metric applied to all Application Monitoring Flows,
- P_L : The duration specified as a positive integer number of milliseconds used for calculating the One-way Packet Loss Ratio Performance Metric applied to all Application Monitoring Flows,
- P_I : The duration specified as a positive integer number of milliseconds used for calculating the Ingress Measured Information Rate Performance Metric applied to all Application Flows,
- P_E : The duration specified as a positive integer number of milliseconds used for calculating the Egress Measured Information Rate Performance Metric applied to all Application Flows, and
- P_T : The duration specified as a positive integer number of milliseconds used for calculating all Performance Metrics applied to all TVCs.



Consequently, there are six sequences of Performance Metric Calculation Intervals with ($k = 0, 1, \dots$).

Note that different Performance Metrics being measured on the same Monitored Entity (other than TVCs) might be calculated over different sequences of Performance Metric Calculation Intervals.

- [R10] The Subscriber and SP **MUST** agree on the value for each of P_D, P_V, P_L, P_I, P_E , and P_T .
- [R11] A value of each Performance Metric **MUST** be calculated for each Performance Metric Calculation Interval in the sequence of Performance Metric Calculation Intervals that applies to the Performance Metric and that is contained in a time interval during which the SD-WAN Service is provided to the Subscriber.
- [R12] The value of a Performance Metric for a Performance Metric Calculation Interval in the sequence of Performance Metric Calculation Intervals that applies to the Performance Metric **MUST** be based on measurements taken during the Performance Metric Calculation Interval.
- [D1] The values of P_D, P_V, P_L, P_I, P_E , and P_T **SHOULD** all be less than or equal to 10000 milliseconds.

7.2.3.3 Performance Metric Calculation Instance

A Performance Metric Calculation Instance is defined as the measurement of a given Performance Metric, using the appropriate Performance Metric Calculation Interval Duration, over a given Monitored Entity. The attributes of a Performance Metric Calculation Instance are defined within this section. The Performance Metric Calculation Instance references the following as defined in Table 9:

- Monitored Entity
- Performance Metric
- Performance Metric Calculation Interval Duration (i.e. the value of P)

Attribute	Value	Value Constraints
Monitored Entity	Identifier for the Monitored Entity	As described in section 7.2.3.1.
Performance Metric	The Performance Metric that is included in the Performance Metric Calculation Instance	For Application Monitoring Flows one of: <ul style="list-style-type: none"> • One-way Mean Packet Delay



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Attribute	Value	Value Constraints
		<ul style="list-style-type: none"> One-way Inter-Packet Delay Variation One-way Packet Loss Ratio For Application Flows one of the following: <ul style="list-style-type: none"> Ingress Measured Information Rate Egress Measured Information Rate For TVCs one of the following: <ul style="list-style-type: none"> One-way Mean Packet Delay One-way Inter-Packet Delay Variation One-way Packet Loss Ratio
Performance Metric Calculation Interval Duration	The duration in milliseconds of the Performance Metric Calculation Interval	A real number from 1-10000 as per [D1].

Table 9 – Performance Metric Calculation Instance Attributes

An example of a Performance Metric Calculation Instance for a TVC would include an Entity Identifier value of *Ingress UCS EP Identifier = 5West, Egress UCS EP Identifier = 5East, UCS CoS Name = Gold*, a Performance Metric *One-way Mean Packet Delay*, and a Performance Metric Calculation Interval Duration of *100 milliseconds*.

[R13] A given Monitored Entity **MUST NOT** appear in more than one Performance Metric Calculation Instance for a given Performance Metric.

The method used to delete a Performance Metric Calculation Instance and stop monitoring, and retrieval of historical Performance Metric values are beyond the scope of this standard.

7.2.4 Performance Metric Definitions

The subsections below define Performance Metrics for Application Monitoring Flows, Application Flows and TVCs.

7.2.4.1 Performance Metrics for Application Monitoring Flows

The definitions for the Performance 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 for a given Ingress UNI for the Ingress Application Flow to an Egress AFMP for 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 Performance 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 whose first bit was received at the Ingress AFMP during the Performance Metric Calculation Interval. Then the One-way Mean Packet Delay for the Application Monitoring Flow over that interval is the arithmetic mean of the values $\delta_1 \dots \delta_n$.
- If Δ is *null* (no measurement values) then the One-way 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 Performance 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 whose first bit was received at the Ingress AFMP during the Performance Metric Calculation Interval.
- Let Δ' = the set of all pairs of elements $\{\kappa_m, \kappa_n\}$ in Δ such that $n > m$ and the difference in the time of transmission at the Ingress AFMP of packets n and m equals the *arrival interval* in the value of SWVC Performance Time Intervals Service Attribute as defined in section 9.8 of MEF 70.1 [10].
- If Δ' is *null* (no measurement values) then the value of the One-way Mean Inter-Packet Delay Variation for the Performance 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 Monitoring Flow over the Performance Metric Calculation Interval is the arithmetic mean of the values v_{mn} .

One-way Packet Loss Ratio for an Application Monitoring Flow and for a Performance Metric Calculation Interval is defined as:

- Let s represent the total number of IP Packets in the Application Monitoring Flow whose first bit was received at the Ingress AFMP during the Performance 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 whose first bit was received at the Ingress AFMP during the Performance Metric Calculation Interval. Duplicated packets are ignored. 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.
 - If $s>0$, then the value of the One-way Packet Loss Ratio is $100*(s-r)/s$

[R14] The One-way Packet Loss Ratio **MUST** be represented as a percentage.

7.2.4.2 Performance Metrics for Application Flows

Measured Information Rate is defined as the Information Rate expressed in bits per second measured during a single Performance Metric Calculation Interval for each Monitored Entity as defined in section 7.2.3.1.

Application Flow Ingress Measured Information Rate in bits per second for a given Ingress Application Flow and a given Performance 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 Performance 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 Performance 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 Performance Metric Calculation Interval Duration for Egress MIR for Application Flows).

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

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

7.2.4.3 Performance Metrics for TVCs

The definitions for the Performance Metrics for TVCs are detailed below. The impact of TVC Performance Metric values on forwarding decisions is beyond the scope of this standard.

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 whose first bit was received at the Ingress TVC MP during the Performance Metric Calculation Interval until the reception of the last bit of the corresponding IP Packet at the Egress TVC MP. If the IP Packet is erroneously duplicated as it traverses the network, the delay is based on the first copy that is delivered.

One-way Mean Packet Delay for packets flowing in one direction over a given TVC from an ingress TVC MP to an egress TVC MP, and for a given Performance 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 whose first bit was transmitted from the ingress TVC MP during the Performance 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 Performance 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 Performance 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 whose first bit was transmitted from the ingress TVC MP during the Performance Metric Calculation Interval.
- Let Δ' = the set of all pairs of elements $\{\kappa_m, \kappa_n\}$ in Δ such that $n > m$ and the difference in the time of transmission at the Ingress AFMP of packets n and m equals the *arrival interval* in the value of SWVC Performance Time Intervals Service Attribute as defined in section 9.8 of MEF 70.1 [10].
- If Δ' is *null*, then the value of the One-way Mean Packet Delay Variation for the Performance 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} .

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 Performance Metric Calculation Interval is defined as:

- Let s represent the total number of IP Packets whose first bit was transmitted over the TVC from Ingress TVC MP to Egress TVC MP during the Performance Metric Calculation Interval.



- Let r represent the total number of IP Packets received from Ingress TVC MP by the Egress TVC MP on the TVC whose first bit was transmitted during the same Performance Metric Calculation Interval. Duplicate packets are ignored. Then the One-way Packet Loss Ratio over that Performance Metric Calculation Interval for the TVC is defined as follows:
 - If $s=0$, then the One-way Packet Loss Ratio is 0.
 - If $s>0$, then the One-way Packet Loss Ratio is $100*(s-r)/s$

[R15] The One-way Packet Loss Ratio **MUST** be expressed as a percentage.

7.3 PM Metric Calculation and Reporting

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

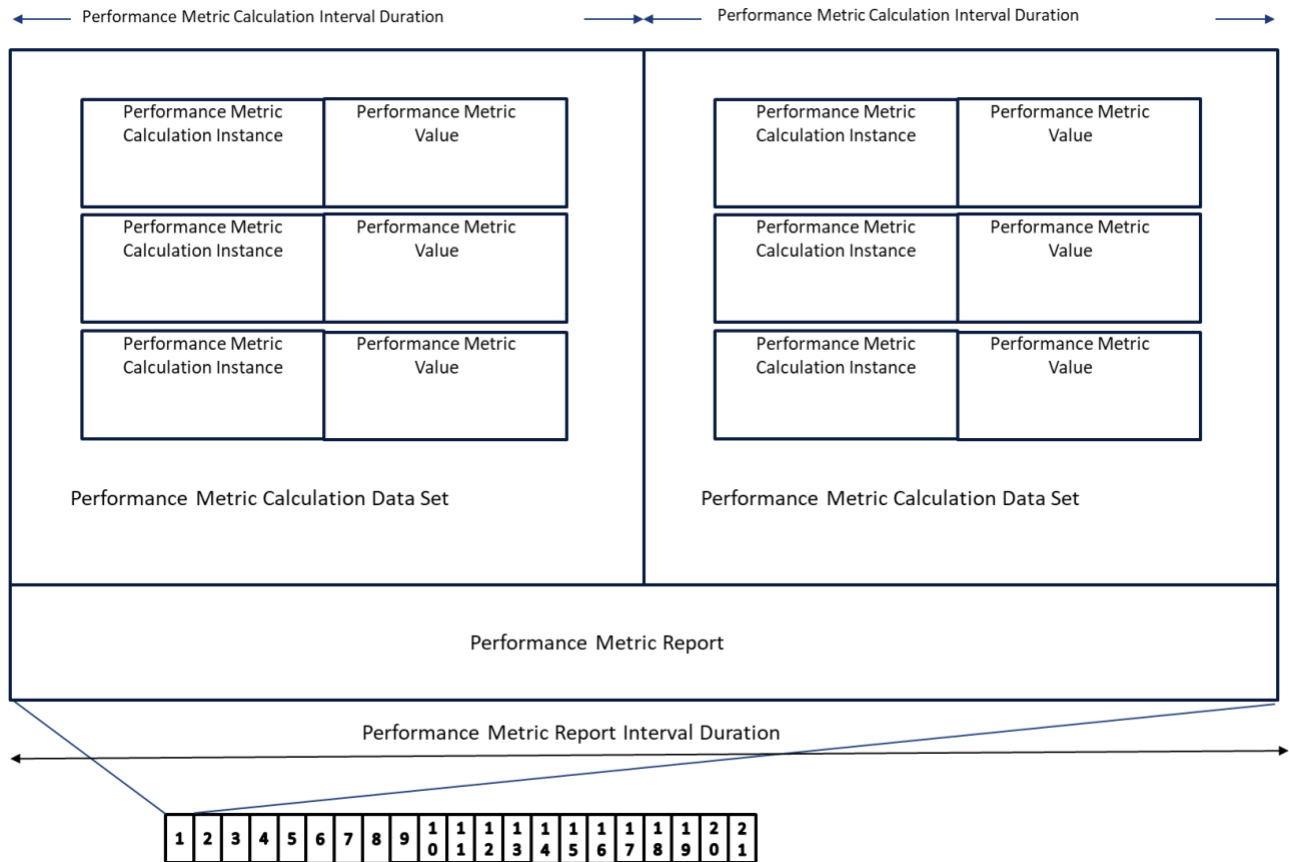


Figure 7 – Performance Metric Calculation and Reporting Example

As shown in Figure 7, in this example a Performance Metric Calculation Data Set, defined as a set of Performance Metric values for a given Performance Metric Calculation Interval, contains three Performance Metric Calculation Instances that are each associated with a Performance Metric Value. A Performance Metric Calculation Interval has a defined duration, the Performance Metric Calculation Interval Duration. Multiple Performance Metric Calculation Instances using the same Performance Metric Calculation Interval Duration and their associated Performance Metric values are combined in a Performance Metric Calculation Data Set. Two Performance Metric Calculation Data Sets are included in a Performance Metric Report. The Performance Metric Report Interval Duration (see text below) defines the length of the report interval. Performance Metric Reports are generated at the regular interval defined by the Performance Metric Report Interval Duration shown as 1 through 21 in the figure. This example only shows one of the six possible sequences of Performance Metric Calculation Intervals (which must therefore be the one for TVCs, using P_T (defined in section 7.2.3.2), since monitoring is always enabled for TVCs). In this example, the Performance Metric Report Interval Duration is twice as long as the Performance Metric Calculation Interval Duration.

The Performance Metric Calculation Data Set is defined as the set of Performance Metric values for a given Performance Metric Calculation Interval, for the Performance Metrics that have been

agreed to be monitored for Application Flows or mandated to be monitored for all TVCs. There is a sequence of Performance Metric Calculation Data Sets for each of the six Performance Metric Calculation Interval Durations described in section 7.2.3.2 and agreed to per [R10]. Within each sequence, there is a Performance Metric Calculation Data Set for each Performance Metric Calculation Interval. A Performance Metric Calculation Data Set contains the following attributes:

- Performance Metric Calculation Data Set Start Time
- Performance Metric Calculation Data Set Duration
- List of pairs each consisting of Performance Metric Calculation Instance and Performance Metric value

These attributes are defined in section 7.3.1.

Reporting of performance data to the Subscriber is based on a sequence of contiguous time intervals 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 an arbitrary UTC time chosen by the Service Provider that is before the SD-WAN service is first provided to the Subscriber, P is a duration, and $k = 0, 1, 2, \dots$. Each interval starting at $t + kP$ and ending at $t + (k+1)P$ is referred to as a Performance Metric Report Interval. The duration P is referred to as the Performance Metric Report Interval Duration. Each Performance Metric Report Interval results in a Performance Metric Report that is reported by the SP to the Subscriber.

Performance Metric Reports contain one or more Performance Metric Data Sets. The number of Performance Metric Data Sets in a Performance Metric Report depends on the Performance Metric Report Interval Duration and the Performance Metric Calculation Interval Duration for the Performance Metric Data Sets.

As an example, a Performance Metric Report has a Performance Metric Report Interval Duration of 5 minutes. The Performance Metric Report contains all Performance Metric Calculation Data Sets that complete during that 5 minute period.

[R16] The Subscriber and SP **MUST** agree on the value of the Performance Metric Report Interval Duration and the Performance Metric Report Start Time.

Performance Metric Reports are generated at the completion of each Performance Metric Report Interval. This process repeats for each Performance Metric Report Interval. The Performance Metric Report contains the following attributes:

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



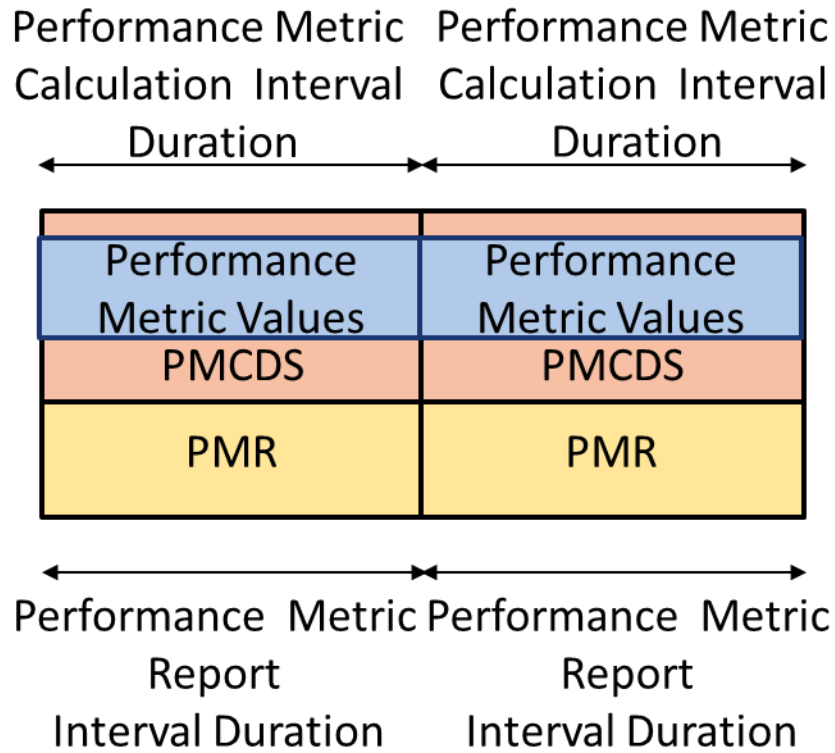
- [R17] The Performance Metric Report Interval Duration **MUST** be an integer multiple of each of the Performance Metric Calculation Interval Durations agreed per [R11].

This ensures that for each Performance Metric Calculation Interval Duration (and corresponding sequence of Performance Metric Calculation Data Sets), every Performance Metric Report contains the same number of Performance Metric Calculation Data Sets from within that sequence.

Note: the Performance Metric Report could include both multiple Performance Metric Calculation Data Sets for the same Performance Metric Calculation Instance but different Performance Metric Calculation Intervals (i.e. covering different periods of time), and Performance Metric Calculation Data Sets for different Performance Metric Calculation Instances (i.e. different Monitored Entities or different Performance Metrics).

The Performance Metric Report, Performance Metric Report Interval, Performance Metric Report Interval Duration, and Performance Metric Calculation Data Set are defined in section 7.3.2.

Some example relationships between Performance Metric Calculation Data Sets and Performance Metric Reports are shown in Figure 8, Figure 9, and Figure 10 below. Not all possible examples are shown in these figures.

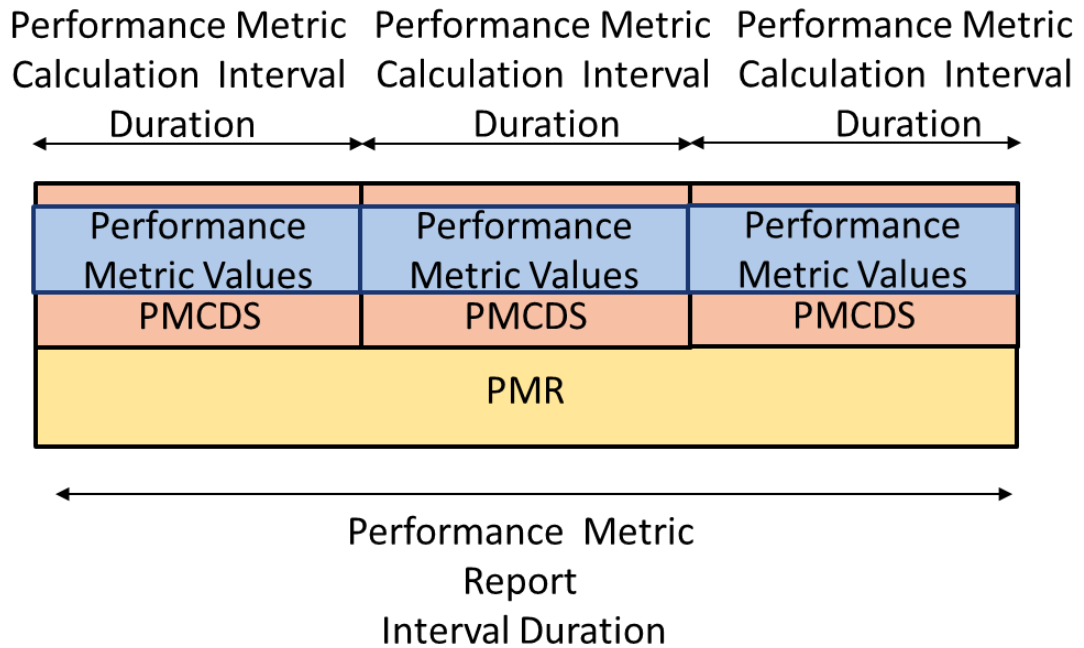


PMCDS– Performance Metric Calculation Data Set
 PMR – Performance Metric Report

Figure 8 – Performance Metric Calculation Interval Duration Equals Performance Metric Report Interval Duration

Figure 8 shows one example of one of the six possible sequences of Performance Metric Calculation Intervals where the value of the Performance Metric Calculation Interval Duration and the value of the Performance Metric Report Interval duration are equal, and the start time of the Performance Metric Calculation Interval and the start time of the Performance Metric Report Interval are aligned.

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



PMCDS– Performance Metric Calculation Data Set

PMR – Performance Metric Report

Figure 9 – Performance Metric Report Interval Greater Than Performance Metric Calculation Interval Duration

In Figure 9 one example of one of the six possible sequences of Performance Metric Calculation Intervals where the value of the Performance Metric Calculation Interval Duration is 1/3 of the Performance Metric Report Interval Duration and the start time of the Performance Metric Calculation Interval Durations are aligned, there are three Performance Metric Calculation Data Sets in each Performance Metric Report. This example illustrates that multiple PM Calculation Data Sets in the same sequence (i.e. for the same PM Calculation Interval Duration) can be included in the same Performance Metric Report.

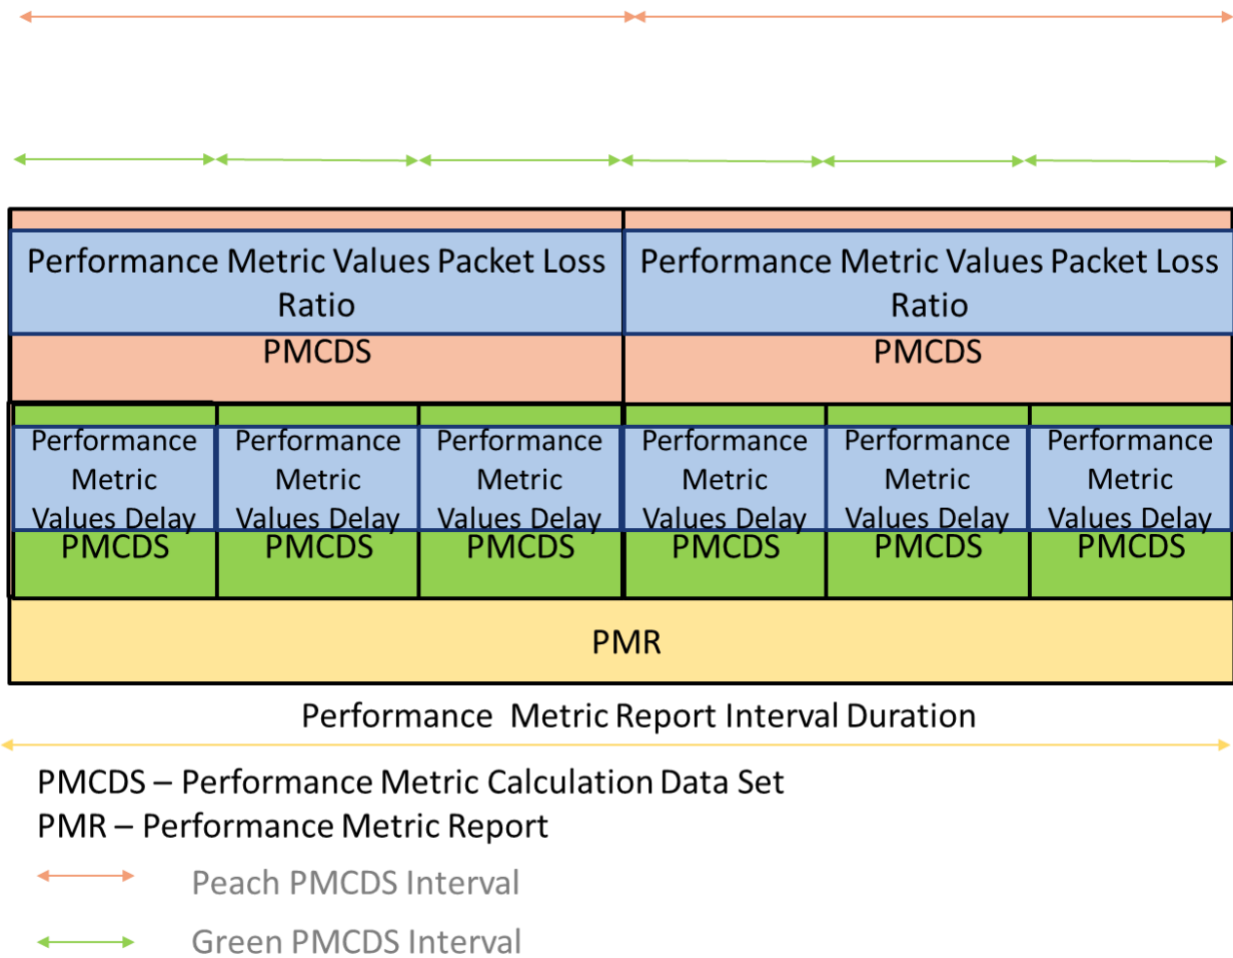


Figure 10 – Different Performance Metric Calculation Interval Durations in a Performance Metric Report

Figure 10 shows an example of when the value of the Performance Metric Calculation Interval Duration is not equal for all Performance Metrics. In this example, there are two different Performance Metric Calculation Interval Durations, shown by peach or green arrows, and associated Performance Metric Calculation Data Sets, shown in peach and green. These Performance Metric Calculation Data Sets are included in the same Performance Metric Report since the Performance Metric Calculation Data Sets are all created within the same Performance Metric Report Interval Duration, shown as the yellow arrow.

The retrieval of historical Performance Metric values is left to the implementation and is beyond the scope of the standard.

7.3.1 Performance Metric Calculation Data Set

The attributes of a Performance Metric Calculation Data Set are defined within this section as shown in Table 10.

Attribute	Value
Performance Metric Calculation Data Set Start Time	The actual Date and Time in UTC that the Performance Metric Calculation Data Set started measurement and calculation for the Performance Metric Calculation Instances in the data set.
Performance Metric Calculation Data Set Duration	The length of time of the Performance Metric Calculation Data Set Duration in ms. This may differ slightly from the intended duration.
Performance Metric Values	A list of three tuples each consisting of : <ul style="list-style-type: none"> • Performance Metric Calculation Instance Monitored Entity as described in Table 9 • Performance Metric as described in Table 9 • The calculated value(s) of the Performance Metric(s) for that Performance Metric Calculation Instance.

Table 10 – Performance Metric Calculation Data Set Attributes

The Performance Metric Calculation Data Set Duration represents the actual length of time of the Performance Metric Calculation Data Set Duration and may differ slightly from the duration agreed to by the Subscriber and SP due to processing time and the actual start time. The bounds on the deviation are beyond the scope of this standard.

Note: all the entries in the list of Performance Metric Values are necessarily for the same type of Monitored Entity; and unless they are all TVCs, all the entries in this list must also be for the same Performance Metric.

[R18] An SD-WAN PM Implementation **MUST** include in the PM Metric Report, the values of the attributes defined in Table 10 for the Performance Metrics that are agreed to be shared with the Subscriber per [R2].

[R19] For each Performance Metric Calculation Instance included in the Performance Metric Calculation Data Set, the Performance Metric Calculation Data Set **MUST** include exactly one Performance Metric value as shown in Table 10.

7.3.2 Performance Metric Report

The attributes of a Performance Metric Report are defined in Table 11.

Attribute	Value
Performance Metric Report Interval Start Time	The start time of the Performance Metric Report interval in UTC.

Attribute	Value
Performance Metric Report Interval Duration	The actual duration, an integer >0 in ms, of the Performance Metric Report.
List of Performance Metric Calculation Data Sets	List of performance Metric Calculation Data Sets

Table 11 – Performance Metric Report Attributes

- [R20]** The SD-WAN PM Implementation **MUST** provide a Performance Metric Report that contains the attributes specified in Table 11 as agreed to per [R2].
- [R21]** An SD-WAN PM Implementation **MUST** generate a Performance Metric Report at the completion of each Performance Metric Report Interval.

Note: there may be some difference in time between the end of the Performance Metric Report Interval 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 Performance Metric Calculation Data Set included in the Performance Metric Report Interval.

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

The period of time agreed to by the SP and the Subscriber might be the subject of legal requirements. Those requirements and the duration of the period of time are outside the scope of this standard.

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

7.4 Threshold Crossing Alert

A Threshold Crossing Alert (TCA) is defined as a notification message that is specific to a particular Performance Metric, a particular set of TCA Function parameters and a particular Monitored Entity. TCAs are generated when the Performance Metric value exceeds, equals, or falls below the TCA Performance Threshold. TCAs can be configured for certain Performance



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Metrics (One-way Mean Packet Delay, One-way Mean Inter-Packet Delay Variation, One-way Packet Loss Ratio) and used to detect when the performance is degraded beyond a given pre-configured value (i.e., the threshold value) for a given Performance Metric. 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 standard the term TCA Function describes the process of generating TCAs for a given Performance Metric Calculation Instance, and a given set of TCA Function parameters values. A TCA Function is the application of a given set of TCA Function parameter values to a given Performance Metric Calculation Instance.

There are two types of TCA reporting, Stateful and Stateless. Stateful TCA reporting is defined as a TCA reporting mechanism whereby a SET-TCA is generated when an alertable condition begins, and a CLEAR-TCA is generated when it ends. It is used to possibly reduce the total number of TCAs that are generated. The intent of Stateful TCA reporting is to provide a SET-TCA notification when a degradation is first encountered, followed by a CLEAR-TCA notification when the degradation is resolved.

Stateless TCA reporting differs from Stateful TCA reporting in that TCAs can be generated whenever a degradation is encountered for a Performance Metric Calculation Interval.

The TCA Performance Threshold is defined as the value against which the Performance Metric value is compared when determining whether or not to generate a TCA.

The TCA Function parameters that are agreed-on by the Subscriber and the SP for each TCA Function are shown in Table 12



Performance Monitoring and Service Readiness Testing for SD-WAN Service

TCA Function Parameters	Description	Possible Values	Comments
Performance Metric Calculation Instance	The information that needs to be agreed is the Monitored Entity identifier (as described in section 7.2.3.1) and the Performance Metric name.	See section 7.2.3.3	
TCA Reporting Type	Indicates whether the TCA reporting is Stateful or Stateless	<i>Stateful</i> or <i>Stateless</i>	
TCA Performance Threshold	The value for the Performance Metric that must be exceeded to trigger a TCA	The three metrics are expressed as One-way Mean Packet Delay (ms), One-way Mean Inter-Packet Delay Variation (ms), and One-way Packet Loss (%)	
SET-TCA Window Threshold	The number of Performance Metric Calculation Intervals, within the TCA Window Size, for which the Performance Metric value must be at or above the TCA Performance Threshold to generate a SET-TCA, when using Stateful TCA Reporting.	Integer >0	Stateful only
CLEAR-TCA Window Threshold	The number of Performance Metric Calculation Intervals, within the TCA Window Size, for which the Performance Metric value must be below the TCA Performance Threshold to generate a CLEAR-TCA, when using Stateful TCA Reporting.	Integer >0	Stateful only
TCA Window Size	The number of consecutive Performance Metric Calculation Intervals in a sliding window that are used to evaluate whether to generate a SET-TCA or CLEAR-TCA, when using Stateful TCA Reporting.	Integer >0	Stateful only
Damping Factor	An integer > 0 that is used to suppress new TCAs.	Integer >0	Stateless only

Table 12 – TCA Function Parameters

[R22] If TCAs are shared with the Subscriber, the values of the TCA Function parameters described in Table 12 **MUST** be agreed-on by the SP and Subscriber for each TCA Function.

[R23] 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.

How reporting of TCAs to the Subscriber is started, stopped, or modified is beyond the scope of this standard.

Note that a TCA does not have to be established for every Performance Metric that is being calculated. Reporting of TCAs is described in section 7.4.3.

Within this standard a single TCA Function is specific to a given Performance Metric Calculation Instance, but it might result in many TCAs. Multiple Performance Metric Calculation Instances for the same Monitored Entity (but for different Performance Metrics) can have TCA Functions associated with them. Similarly, each Performance Metric Calculation Instance can have multiple TCA Functions associated with it (e.g., with different TCA Parameter values).

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

The TCA Function using Stateful TCA Reporting has an internal state, whose value is 'set' or 'clear'. The TCA Function algorithm is informally described in the following two paragraphs and the flowchart in Figure 11. The algorithm is defined in the requirements in section 7.4.3.2.

The TCA Function initializes in the 'clear' state and stays in the 'clear' state for the first Performance Metric Calculation Interval and all subsequent Performance Metric Calculation Intervals until such time that the Performance Metric value has exceeded the Threshold. A SET-TCA is generated when the Performance Metric value is equal to or greater than the TCA Performance Threshold value for the number of Performance Metric Calculation Intervals defined by the value of the SET-TCA Window Threshold out of the number of Performance Metric Calculation Intervals defined by TCA Window Size. The TCA Function is then considered to be in a 'set' state, and no further SET-TCAs are generated while in this state.

The TCA Function moves from the 'set' state to the 'clear' state when the Performance Metric value is less than the TCA Performance Threshold value for the number of Performance Metric Calculation Intervals defined by the value of the CLEAR-TCA Window Threshold out of the number of Performance 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 11. The flowchart does not specify a specific implementation of the algorithm. The requirements contained in section 7.4.3 define the algorithm and are controlling in case of any conflict between the requirements and the flowchart below. 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 $j = n, n - 1, n - 2, \dots, n - W + 1$

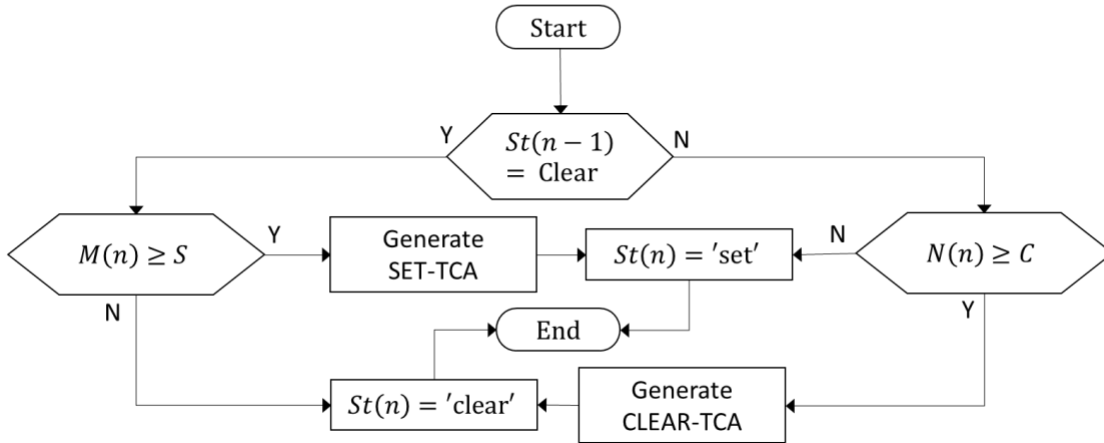


Figure 11 – Flowchart for Stateful TCA

The algorithm is formally defined in requirements [CR11]< and [CR12]< in section 7.4.3.2.

Figure 12 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 [CR8]< in section 7.4.3.2). This is mandated to avoid a condition where the SET-TCA and CLEAR-TCA criteria are met at the same time.

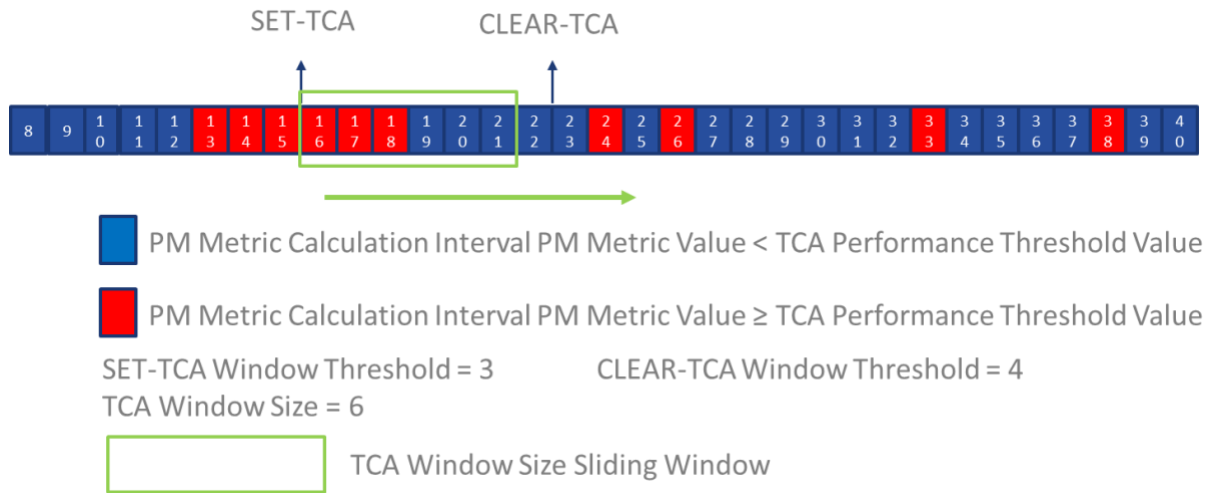


Figure 12 – Stateful TCA with Consecutive Performance Metric Calculation Intervals ≥ TCA Performance Threshold Value Example

In the example in Figure 12, the TCA Window Size is set to 6 and moves to the right in the figure with the completion of each Performance 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 Performance Metric value is equal to or greater than TCA Performance Threshold value for SET Threshold value in Performance 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 Performance Metric value is less than TCA Performance Threshold value for CLEAR Threshold value in Performance 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 Performance Metric value is equal to or greater than the TCA Performance Threshold value in Performance 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 13 – Stateful TCA without Consecutive Performance Metric Calculation Intervals ≥ TCA Performance Threshold Value Example

Figure 13 shows an example of a SET-TCA and CLEAR-TCA being generated even if consecutive Performance Metric Calculation Intervals do not meet the criteria for SET-TCA or CLEAR-TCA. The Performance Metric value is equal to or greater than the TCA Performance Threshold value in Performance 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 Performance Metric Calculation Interval 16. A CLEAR-TCA is generated after the completion of Performance Metric Calculation Interval 29 because the CLEAR-TCA Window Threshold and TCA Window Size criteria have been met (Performance Metric Calculation Intervals 25, 27-29), even though the Performance Metric value in Performance 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 Performance Metric Calculation Interval separately. To prevent many repeating TCAs, a Damping Factor is defined within this standard.

The algorithm is informally described in the following paragraph and the flowchart in Figure 14. The algorithm is defined in the requirements [CR15] <, [CR17] <, [CR18] <, [CR19] <, [CR20] <, [CR21] <, and [CR22] < in section 7.4.3.3.

When using Stateless TCA reporting, defined as a TCA reporting mechanism that generates a TCA when an alertable condition is detected, each TCA Function has a single configured TCA Performance Threshold value. When a Performance Metric value in a Performance Metric Calculation Interval is equal to or greater than the TCA Performance Threshold value for a Performance Metric Calculation Interval, a TCA can be generated. There is no corresponding CLEAR.

The Damping Factor is defined as an integer > 0 that is used to suppress new TCAs. The Damping Factor is always present and does not suppress any TCAs when it is defined as 1. The Damping

Factor defines the number of Performance Metric Calculation Intervals for which new TCAs are suppressed after a TCA is sent.

The algorithm for Stateless TCAs with the Damping Factor value = 1 is informally shown in Figure 14.

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

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




Figure 14 – Flowchart for Stateless TCA when Damping Factor Value = 1

Figure 15 shows an example of a Stateless TCA with the Damping Factor equal to 1.

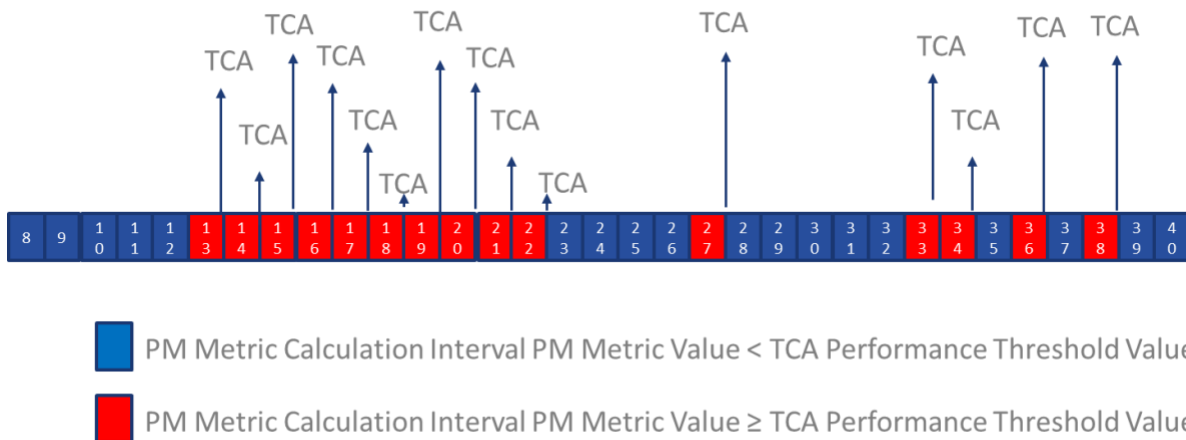


Figure 15 – Stateless TCA with Damping Factor = 1 Example

As shown in Figure 15, a TCA is generated for each Performance Metric Calculation Interval in which a degraded condition exists. In this example, the TCA criterion is met in Performance

Metric Calculation Interval 13-22, 27, 33, 34, 36, and 38. TCAs can be generated for each of these Performance Metric Calculation Intervals.

The Damping Factor is used to avoid generating a TCA for every Performance Metric Calculation Interval with a degraded condition.

The algorithm for Stateless TCA with Damping Factor is informally described in Figure 16. Note that the flowchart in Figure 16 is equivalent to the flowchart in Figure 14 when the Damping Factor value = 1. 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, [CR15] <, [CR17] <, [CR18] <, [CR19] <, [CR20] <, [CR21] <, and [CR22] < in section 7.4.3.3.

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

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

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

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

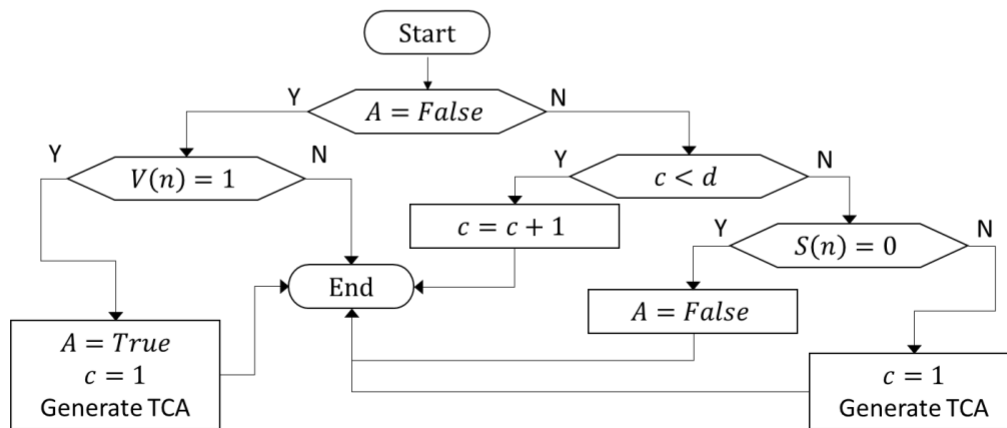


Figure 16 – Flowchart for Stateless TCA w/Damping Factor ≥ 2

The impact of the Damping Factor is shown in Figure 17.

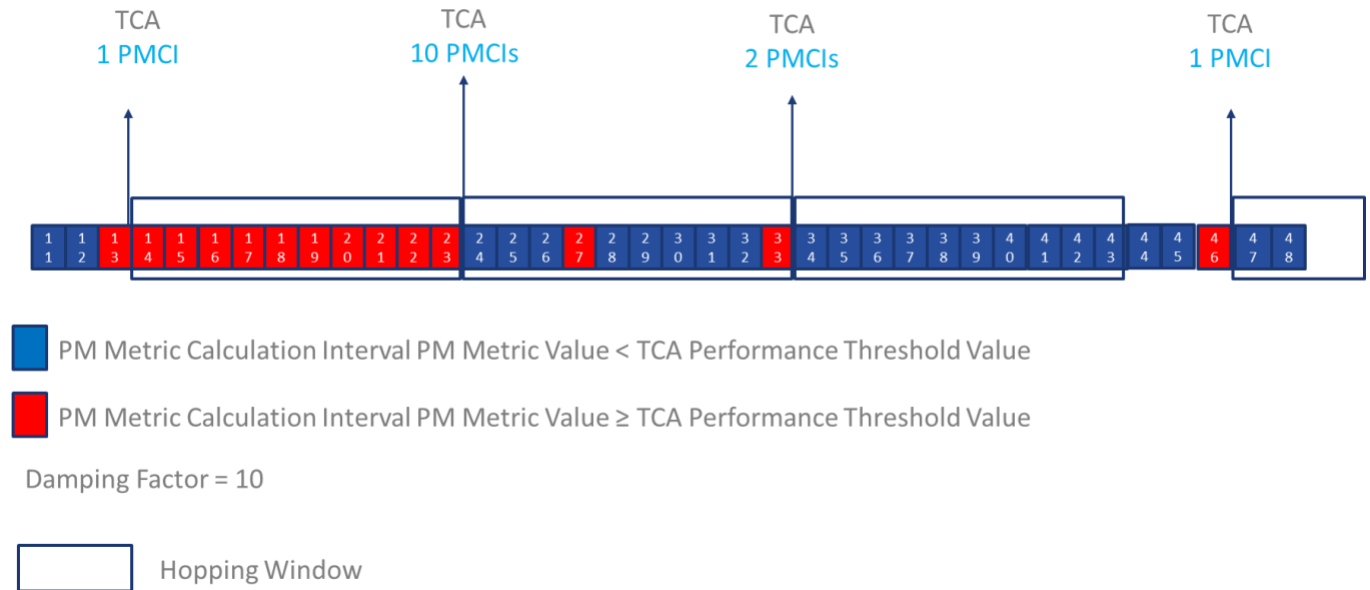


Figure 17 – Stateless TCA with Damping Factor = 10 Example

As shown in Figure 17, rather than multiple TCAs being generated, the TCAs are suppressed for the number of consecutive Performance Metric Calculation Intervals defined by the Damping Factor value. In the example, the Performance Metric value is equal to or greater than the TCA Performance Threshold value in Performance Metric Calculation Interval 13. An initial TCA is generated at the end of the Performance Metric Calculation Interval, and the Damping algorithm begins suppressing TCAs. A hopping window that is equal to the number of Performance Metric Calculation Intervals specified by the Damping Factor value starts at the completion of Performance Metric Calculation Interval 13.

Moving forward 10 Performance Metric Calculation Intervals (Damping Factor value) another TCA is generated at the completion of Performance Metric Calculation Interval 23 since at least one Performance Metric Calculation Interval within the hopping window had a Performance Metric value equal to or greater than the TCA Performance Threshold value. In this example, this TCA includes the count of Performance Metric Calculation Intervals in the hopping window that had a Performance Metric value equal to or greater than the TCA Performance Threshold value (shown as 10 PMCIs in blue text in Figure 17); the inclusion of the number of Performance Metric Calculation Intervals in the Damping Factor TCA Notification Message Field is a recommended parameter per [CD3]<.

Moving forward another 10 Performance Metric Calculation Intervals to Performance Metric Calculation Interval 33, a TCA is generated since at least one Performance Metric Calculation Interval within the hopping window had a Performance Metric value equal to or greater than the TCA Performance Threshold value. This TCA includes the count of Performance Metric Calculation Intervals in the hopping window that had a Performance Metric value that is equal to or greater than the TCA Performance Threshold value (shown as 2 PMCIs in blue text in Figure 17). Moving forward 10 more Performance Metric Calculation Intervals to interval 43, a TCA is



not generated since zero of the Performance Metric Calculation Intervals in the hopping window had a Performance Metric value is equal to or greater than the TCA Performance Threshold value.

At the completion of a hopping window without the occurrence of any TCAs, the Damping Factor Algorithm stops suppressing TCAs. If a future Performance Metric Calculation Interval has a Performance Metric value equal to or greater than the TCA Performance Threshold value, a TCA is generated at the completion of the Performance Metric Calculation Interval, and the Damping Factor Algorithm is activated.

This behavior repeats as long as a Stateless TCA occurs during the hopping window which is shown by the TCA at the completion of Performance Metric Calculation Interval 46 which starts suppressing TCAs 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 TCA Requirements that apply to both Stateful and Stateless, requirements that only apply to Stateful TCAs, and requirements 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.

[D3] An SD-WAN PM Implementation **SHOULD** provide the ability to generate TCAs.

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

[CR2]<[D3] An implementation supporting TCAs **MUST** support TCA Functions with different TCA Function parameter values for each Performance Metric Calculation Instance.

[CD1]<[D3] An implementation supporting TCAs **SHOULD** support multiple TCA Functions with different TCA Function parameter values for a given Performance Metric Calculation Instance.

[CR3]<[D3] If a TCA is generated for a given Performance Metric Calculation Interval, then it **MUST** be generated by the end of the Performance Metric Calculation Interval following the one that caused the TCA to be generated.

The SD-WAN SP can report the occurrence of a TCA to the Subscriber as defined in [R23].

Note: the method used to convey TCAs to the Subscriber is beyond the scope of this standard.

7.4.3.2 Stateful TCA Requirements

The requirements in this section apply to Stateful TCA implementations.



Performance Monitoring and Service Readiness Testing for SD-WAN Service

- [CR4]<[D3] An implementation supporting Stateful TCA Reporting **MUST** support all the TCA Function parameters listed in Table 12 except those indicated as being for Stateless only.
- [CR5]<[D3] An implementation supporting Stateful TCA Reporting **MUST** be able to configure the value of the SET-TCA Window Threshold being any integer value within a range of 1-300.
- [CR6]<[D3] An implementation supporting Stateful TCA Reporting **MUST** be able to configure the value of the CLEAR-TCA Window Threshold being any integer value within a range of 1-300.
- [CR7]<[D3] An implementation supporting Stateful TCA Reporting **MUST** be able to configure the value of the TCA Window Size being any integer value within a range of 1-300.

The value of 300 is based on an assumption that Performance Metric Calculations are occurring once per second. Network reconvergence within the SD-WAN service is assumed to have taken place within 300 seconds.

- [CR8]<[D3] 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 strictly greater than the value of the TCA Window Size.
- [CR9]<[D3] For an implementation supporting Stateful TCA Reporting, for each TCA Function, the value of the SET-TCA Window Threshold **MUST** be less than or equal to the value of the TCA Window Size.
- [CR10]<[D3] For an implementation supporting Stateful TCA Reporting, for each TCA Function, the value of the CLEAR-TCA Window Threshold **MUST** be less than or equal to the value of the TCA Window Size.
- [CR11]<[D3] An implementation supporting Stateful TCA Reporting **MUST** generate a SET-TCA notification message at the end of a Performance 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 Performance Metric value equals or exceeds the Performance Threshold in at least S of the last W Performance Metric Calculation Intervals (including the one just ended), where S is the SET-TCA Window Threshold and W is the TCA Window Size.
- [CR12]<[D3] An implementation supporting Stateful TCA Reporting **MUST** generate a CLEAR-TCA notification message at the end of a Performance 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 Performance Metric value is less than the Performance Threshold in at least C of the last W Performance



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Metric Calculation Intervals (including the one just ended), where C is the CLEAR-TCA Window Threshold and W is the TCA Window Size.

[CR13]<[D3] An implementation supporting Stateful TCA Reporting **MUST NOT** generate a SET-TCA or a CLEAR-TCA notification message other than as specified in [CR11]< and [CR12]<.

[CR14]<[D3] An implementation supporting Stateful TCA Reporting **MUST** include the attributes shown in Table 13 in the SET-TCA or CLEAR-TCA notification message.



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Field Name	Field Value	Field Description
Monitored Entity	The Monitored Entity identifier	As described in section 7.2.3.1
Date and Time	Date/time in UTC	Time of the event, in UTC. For the SET-TCA and CLEAR-TCA this is the time of the completion of the Performance Metric Calculation Interval for which the Performance Metric value triggered the TCA to be generated.
Performance 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 Metric for which the TCA Function was configured
Performance Metric Values	The Numeric values of the Performance Metric for all the Performance Metric Calculation Intervals within the TCA Window Threshold	Performance Metric values are reflected in milliseconds 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 a Real Number > 0	The TCA Performance Threshold value for the Performance Metric. Milliseconds are used for delay and percentage is used for loss.
TCA Window Threshold	Numeric value as an Integer ≥ 1	For a "STATEFUL-SET" the value is the SET-TCA Window Threshold or for a "STATEFUL-CLEAR" the value is the CLEAR-TCA Window Threshold.
TCA Window Size	Numeric value as an Integer ≥ 1	The value of the TCA Window Size for the SET-TCA or CLEAR-TCA process.
TCA Message Type	<i>STATEFUL-SET</i> , or <i>STATEFUL-CLEAR</i> as a String	The type of TCA, i.e., STATEFUL-SET or STATEFUL-CLEAR
Severity Level	<i>CRITICAL</i> , <i>MAJOR</i> , <i>MINOR</i> , <i>WARNING</i> , or <i>CLEARED</i> 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 13 – Stateful TCA Notification Message Fields

- [R24] If TCAs are shared with the Subscriber, the criteria for each possible value of the Severity Level, specified in Table 13 and Table 14 **MUST** be agreed-on by the SP and Subscriber.

Note: CLEARED is included in the Severity Level to align with TCA implementations where the STATEFUL-CLEAR is communicated by using the CLEARED Severity Level.

7.4.3.3 *Stateless TCA Reporting Requirements*

The requirements in this section apply to Stateless TCA implementations.

- [CR15]<[D3] An implementation supporting Stateless TCA Reporting **MUST** support all the TCA Function parameters listed in Table 12 except those indicated as being for Stateful only.
- [CR16]<[D3] An implementation of Stateless TCA Reporting **MUST** support Damping Factor = 1
- [CR17]<[D3] An implementation of Stateless TCA Reporting with the Damping Factor = 1 **MUST** generate a TCA at the end of a Performance Metric Calculation Interval if, and only if, the Performance Metric value for that Performance Metric Calculation Interval is greater than or equal to the TCA Performance Threshold.
- [CD2]<[D3] An implementation of Stateless TCA Reporting **SHOULD** support the Damping Factor values ≥ 2 .
- [CR18]<[CD2]< An implementation of Stateless TCA Reporting that supports Damping Factor values ≥ 2 **MUST** support Damping Factor integer values between 1-300 Performance Metric Calculation Intervals.
- [CR19]<[CD2]<An implementation of Stateless TCA Reporting with the Damping Factor ≥ 2 **MUST** generate a TCA at the end of a Performance Metric Calculation Interval if the Performance Metric value for that Performance 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 Performance Metric Calculation Intervals, where d is the value of the Damping Factor.
- [CR20]<[CD2]< An implementation of Stateless TCA Reporting with the Damping Factor ≥ 2 **MUST NOT** generate a TCA at the end of a Performance Metric Calculation Interval if a TCA has been generated after any of the $d-1$ preceding Performance Metric Calculation Intervals, where d is the value of the Damping Factor.
- [CR21]<[CD2]< An implementation of Stateless TCA Reporting with the Damping Factor ≥ 2 **MUST** generate a TCA at the end of a Performance Metric Calculation

Interval if a TCA was generated at the end of the Performance Metric Calculation Interval that precedes this one by d Performance Metric Calculation Intervals, and for at least one of the $d-1$ intervening Performance Metric Calculation Intervals, the Performance Metric value was greater than or equal to the TCA Performance Threshold, where d is the value of the Damping Factor.

[CR22]<[CD2]< An implementation of Stateless TCA Reporting with the Damping Factor ≥ 2 **MUST NOT** generate a TCA at the end of a Performance Metric Calculation Interval if the Performance Metric value was less than the TCA Performance Threshold for this Performance Metric Calculation Interval and for all of the preceding $d-1$ Performance Metric Calculation Intervals, where d is the value of the Damping Factor.

[CD3]<[CD2]< For an implementation of Stateless TCA Reporting with the Damping Factor ≥ 2 , when a TCA is generated at the end of a Performance Metric Calculation Interval, it **SHOULD** include the number of Performance Metric Calculation Intervals, within the sequence of d Performance Metric Calculation Intervals ending with the one at the end of which the TCA was generated, in which the Performance Metric value was greater than or equal to the TCA Performance Threshold, where d is the value of the Damping Factor.

Note that the number of Performance Metric Calculation Intervals are included in the first TCA.

[CR23]<[D3] An implementation of Stateless TCA Reporting **MUST** include the information shown in Table 14 in the TCA.

Field Name	Field Value	Field Description
Monitored Entity	The Monitored Entity identifier	As described in section 7.2.3.1
Date and Time	Date/time in UTC	Time of the event, in UTC. This is the time of the end of the Performance Metric Calculation Interval for which the TCA is generated.
Performance 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 Metric for which the TCA Function was configured.
TCA Performance Threshold	Numeric value as a Real Number > 0	The TCA Performance Threshold value for the Performance Metric. Milliseconds are used for delay and percent is used for loss.
Performance Metric Value	The Numeric values of the Performance Metric	The Performance Metric value for the latest Performance Metric Calculation Interval in which the TCA Performance Threshold value was equal to or exceeded by the Performance Metric value.
TCA Message 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 14 – Stateless TCA Reporting Notification Message Fields

[CR24]<[CD2]< If an implementation of Stateless TCA Reporting has the Damping Factor ≥ 2 , the Damping Factor field shown in Table 15 **MUST** be appended to the TCA notification message.

[CR25]<[D3] If an implementation of Stateless TCA Reporting includes reporting the number of Performance Metric Calculation Intervals per [CD3]< [CD2]< , that had a Performance Metric value that is equal to or greater than the TCA Threshold value, the Number of Performance Metric Calculation Intervals shown in Table 15 **MUST** be appended to the TCA Notification.



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

Table 15 – Damping Factor TCA Notification Message Field

There may be questions about how many of the Performance 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 Performance Metric Calculation Intervals exceeding the Performance Threshold Value attribute in Table 15 identifies the number of Performance Metric Calculation Intervals that met the criteria for a TCA. An implementation of the Damping Factor without this desirable capability may cause a Subscriber and/or the SD-WAN SP to perform additional research to determine the number of Performance Metric Calculation Intervals within the hopping window that met the criteria to generate a TCA. This is only recommended since a Performance implementation may not be able to calculate the number of intervals that are impacted after the interval closes.

8 Service Readiness Testing for an SD-WAN Service

Service Readiness Testing (SRT) is the process of testing connectivity across UCSs that exist 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 whether there is connectivity between 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. Verify in this context is used to describe the testing and the reporting of the results.

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 standard. 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)
- Reporting UCS Service Attributes (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 pairs of UCS End Points Connectivity Use Case

This use case describes the verification of ordered pairs 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 whether 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 18 – SRT Order Pairs of UCS End Points Connectivity Verification Use Case Example

Figure 18 shows the use case for the verification of connectivity between SD-WAN Edges. SRMPs (as described in section 8.2) contain a Generator Test Function (GTF) defined as a function for generating and transmitting IP Test Packets and a Collector Test Function (CTF) defined as a function for receiving, counting and discarding received IP Test Packets. The CTF and GTF are each connected to a UCS End Point. The ordered pairs of GTF and CTF are used to perform testing on the ordered pairs of UCS End Points. When a new SD-WAN Service is being activated, tests are performed on all agreed-on ordered pairs of UCS End Points.

When UCS End Points are connected to the Internet, IP reachability is required between the UCS End Points and therefore between the GTFs and CTFs. When both UCS End Points are a UCS Type of *Public* the UCS End Points are assigned either a public or private IP address.

If both UCS End Points are using private IP addresses, then an external coordinator is required to perform this testing. The external coordinator and their functions are outside the scope of this standard.

If One UCS End Point is using a private IP address and the other UCS End Point is using a public IP address, Test IP Packets are generated at the UCS End Point with the private IP address. Test IP Packets are reflected back from the Public end SRMP to the Private end SRMP. If both UCS

End Points are using public IP addresses, Test IP Packets can be generated at the SRMP at either UCS End Point.

When testing across UCSs with a UCS Type of *Private*, the address type assigned to the UCS End Point does not impact the testing.

At least one GTF initiates the testing between the ordered pairs of GTFs and CTFs. The near-end GTF sends test packets to the ordered pair CTF. The far-end GTF either reflects packets received by the far-end CTF back towards the near-end CTF, or generates new packets and transmits them towards the near-end CTF.

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 undergoing SRT. The method 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 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. 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 always receives and counts the IP Test Packets coming from the GTF. In addition, it can either discard or process such IP Test Packets e.g., reflect the IP Test Packets back to the GTF. Testing is performed using only Unicast IP Test Packets, since a GTF is paired with a CTF so that the IP Test Packets generated by the GTF are collected by a particular CTF.

8.3 Service Readiness Measurement Point Locations

The location of SRMPs within the network is defined in this section. Figure 19 shows the location of SRMPs in relationship to processing functions within the SD-WAN Edge. How these functions are implemented, and which tools are used to generate and receive IP Packets are beyond the scope of this standard.

- [R25] An SRMP **MUST** contain a Generator Test Function (GTF) and a Collector Test Function (CTF).
- [R26] The SRMP **MUST** be located so that IP Test Packets generated with the proper format by the GTF are inserted at the UCS UNI which bounds the UCS under test without being processed by the SWVC End Point.
- [R27] The SRMP **MUST** be located so that IP Test Packets collected by the CTF are received at the UCS UNI which bounds the UCS under test without being processed by the SWVC End Point.

- [R28] If any TVCs exist at the time that SRT is performed, the SRMP **MUST** be located so that it connects to the UCS UNI and not to the TVC(s).
- [R29] IP Test Packets **MUST** be formatted so that they are mapped to the appropriate UCS End Point within a UCS UNI.

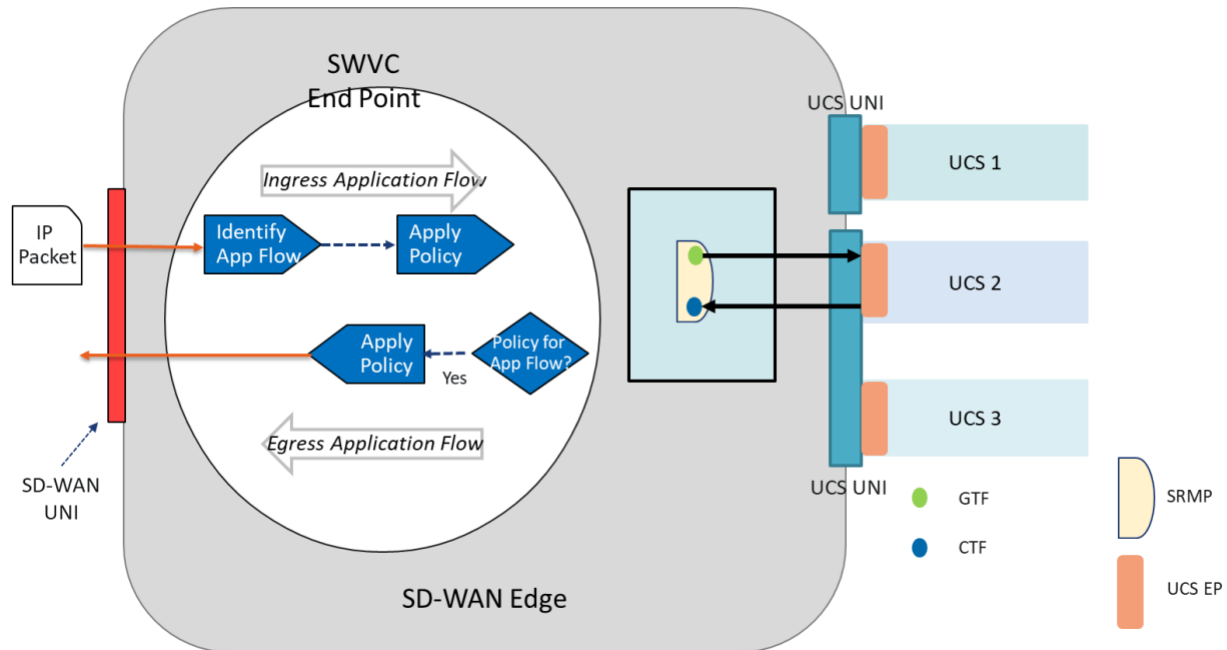


Figure 19 – SRMP Location of UCS End Points Connectivity Testing

8.4 UCS Service Attribute Reporting

This section of the standard 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.

- [R30] The values of UCS Service Attributes (MEF 70.1 [10] section 12), UCS UNI Service Attributes (MEF 70.1 [10] 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 [R30], 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

This section defines the SRT parameters and results for tests that are agreed between the SP and Subscriber. A prerequisite to this testing is verifying that the SD-WAN Controller/Orchestrator can communicate with SD-WAN Edges. How the connectivity with the SD-WAN Controller/Orchestrator is established is beyond the scope of this standard.

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 the function of SRT.

The parameters for UNI Ordered End Point SRT are contained in a 4-tuple of the form $\langle OPairlist, T_{sc}, Testn, Lostn \rangle$ where:

- *OPairlist* is a list of 2-tuples of the form $\langle EPid1, EPid2 \rangle$ where
 - *EPid1* is the value of the UCS End Point Identifier at the UCS UNI where a GTF is located
 - *EPid2* is the value of the UCS End Point Identifier at the UCS UNI where the corresponding CTF is located
- *T_{sc}* is the duration of the test for all the entries in *OPairlist*
- *Testn* is the number of IP packets generated uniformly distributed over *T_{sc}* by the GTF for all the entries in *OPairlist* during the test
- *Lostn* is the maximum number of lost IP packets for the test result to be *pass* for all entries in *OPairlist*

[R31] The SD-WAN SP and the Subscriber **MUST** agree on the values of the parameters of Service Readiness Testing.

The Service Provider and Subscriber can agree on multiple instances of UNI Ordered End Point SRT.

8.6 Test Methodology

This section defines the SRT methodology for verifying connectivity between UCS UNI Ordered End Point Pairs, as depicted in Figure 20.

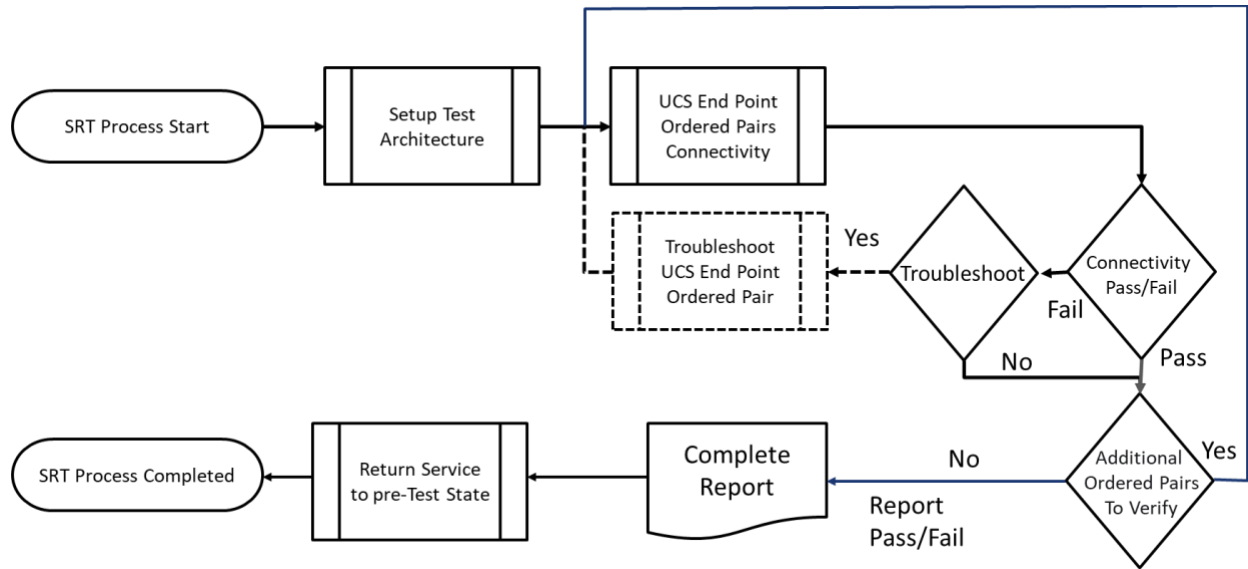


Figure 20 – UCS End Point Ordered Pairs SRT Process Flow

The first step of the process flow is to create the required SRMPs at the appropriate SD-WAN Edges and to connect each SRMP to the correct UCS UNI see Set Up Test Architecture in Figure 20. Depending on the tools and SRT functions used, configurations within the SD-WAN Edge may need to be modified. Any such modifications are beyond the scope of the standard.

The second step in the process flow is to verify connectivity between SRMPs as described in this section. This step is repeated for each ordered pair that is agreed-on to be included in SRT by the Subscriber and SD-WAN SP (see UCS End Point Ordered Pair Connectivity in Figure 20).

The third step in the process flow is to report both the UCS Service Attributes and the UCS UNI Order Endpoint Pair Connectivity Test results (see Complete Report step in Figure 20).

The fourth and final step in the process flow is to disconnect the SRMP from the UCS UNIs and to restore the SD-WAN Edge to its normal configuration (see Return Service to pre-Test State in Figure 20).

Troubleshooting of a test failure is optional and is shown by dashed lines in Figure 20. 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 by the SP, the UCS End Point ordered pair test is repeated when troubleshooting is completed. If troubleshooting is not performed or troubleshooting failed to resolve the problem, the failure is reported.

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 16.

[R32] When a new SWVC is being activated, the SD-WAN SP **MUST** perform SRT between the agreed upon ordered pairs of UCS End Points.



Performance Monitoring and Service Readiness Testing for SD-WAN Service

- [R33] When a new SD-WAN Edge is added to an existing SWVC, the SD-WAN SP **MUST** perform SRT between the Subscriber and SD-WAN SP agreed upon new ordered pairs of UCS End Points.
- [R34] When the SD-WAN SP is performing SRT, the Test Methodology defined in Table 16 **MUST** be used.
- [R35] Results for each ordered pair of UCS End Points tested in the SRT Test Methodology **MUST** be reported as *pass* or *fail*.



Performance Monitoring and Service Readiness Testing for SD-WAN Service

Service Readiness Test Methodology	
Test Name	UCS UNI End Point Ordered Pair SRT
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 the pair of UCS End Points under test, <EP1, EP2>, SRMP1 is placed in the SD-WAN Edge where EP1 is located and SRMP2 is placed in the SD-WAN Edge where EP2 is located. Placement of each SRMP within an SD-WAN Edge is shown in Figure 19. IP reachability is required from the GTF in SRMP1 to the CTF in SRMP2. •SRMP1 GTF offers <i>Testn</i> IP Test Packets with the IP Destination Address for reaching SRMP2 CTF so that the IP Test Packets are injected into the UCS End Point by the SRMP1 GTF at an interval where the <i>Testn</i> 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 CTF 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 calculated lost packets is compared to the <i>Lostn</i> value to determine whether the test is declared pass or fail. •The above is repeated for each ordered pair of UCS End Points agreed to be tested by the Subscriber and SD-WAN SP.
Parameters	$T_{sc}, Testn, Lostn, OPairlist$
Results	<p>Pass = For each entry in <i>OPairlist</i>, the result is Pass if the number of lost packets for that entry in <i>OPairlist</i> is less than or equal to the value of <i>Lostn</i>.</p> <p>Fail = For each entry in <i>OPairlist</i>, the result is Fail if the number of lost packets for that entry in <i>OPairlist</i> is greater than the value of <i>Lostn</i>.</p>

Table 16 – UCS End Point Ordered Pair Connectivity Test Methodology

Note: To test a pair of UCS End Points in each direction, both orders of the pair (A-Z and Z-A) need to be included in the agreed list of ordered pairs of UCS End Points to test.

When the agreed set of ordered pairs of UCS End Points have a Pass result or with a Fail result that is not the SD-WAN SP’s responsibility, the SP is ready for the Subscriber to begin using 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 precise format of the SRT record is beyond the scope of this standard.

[D4] The SRT record **SHOULD** be shared with the Subscriber.

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 10.4, *Subscriber Ethernet Service Attributes*, December 2018
- [5] MEF 35.1, *Service OAM Performance Monitoring Implementation Agreement*, May 2015
- [6] MEF 48.1, *Carrier Ethernet Service Activation Testing*, November 2019
- [7] MEF 61.1, *IP Service Attributes*, May 2019
- [8] MEF 66, *Service OAM for IP Services*, July 2020
- [9] MEF 67, *Service Activation Testing for IP Services*, December 2020
- [10] MEF 70.1, *SD-WAN Service Attributes and Service Framework*, November 2021



Appendix A Service Readiness Testing vs Service Activation Testing (Informative)

This standard 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 determines if 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. An 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. Testing for 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 determines if the agreed to UCSs are ready for the Subscriber to use. The SD-WAN SP is now able to implement TVCs or Policies. SRT does not test the operation of the Policies or the values of the 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.



Appendix B Acknowledgements (Informative)

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

- David **BALL**
- Mike **BENCHECK**
- Neil **DANILOWICZ**
- Bob **KLESSIG**
- Philippe **LOCHON**
- Basil **NAJEM**
- Jessy **ROUYER**
- Ralph **SANTITORO**
- Kamil **SWACZYNA**
- Jason **WOLFE**