



MEF Standard

MEF 70.2

SD-WAN Service Attributes and Service Framework

October 2023

Disclaimer

© MEF Forum 2023. All Rights Reserved.

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

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

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

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

Implementation or use of specific MEF standards, specifications, or recommendations will be voluntary. This document is provided "as is" with no warranties whatsoever, express or implied, including without limitation, any warranties of merchantability, non-infringement, accuracy, completeness or fitness for any particular purpose. MEF and its members disclaim all liability, including liability for infringement of any proprietary rights, relating to use of information in this document.

Table of Contents

1	List of Contributing Members	1
2	Abstract	2
3	Terminology and Abbreviations	3
4	Compliance Levels	7
5	Document Conventions	8
5.1	Numerical Prefix Conventions	8
5.2	Notational Conventions	8
5.3	Diagram Conventions	8
6	Introduction	10
6.1	Document Scope.....	10
6.2	Characteristics of an SD-WAN Service	11
6.3	Organization of the Standard.....	11
7	Key Concepts and Definitions	13
7.1	Service Attributes	13
7.2	SD-WAN Subscriber and SD-WAN Service Provider.....	14
7.3	Subscriber Network and Service Provider Network.....	14
7.4	SD-WAN UNI.....	15
7.5	SWVC End Point.....	16
7.6	SWVC and SD-WAN Service Components Overview	16
7.7	Underlay Connectivity Service, UCS UNI, and UCS End Point	17
7.8	SD-WAN Edge.....	19
7.9	SD-WAN Services Framework	20
7.10	Tunnel Virtual Connection (TVC)	21
7.11	Application Flow, Application Flow Specification, and Policy	23
7.12	Zone	24
7.13	Virtual Topology	25
7.14	Internet Breakout	26
7.15	Rate Limiter and Bandwidth Flow	27
7.16	SD-WAN IP Packet Delivery	31
7.16.1	IP Packet Forwarding.....	31
7.16.2	IP Packet Transparency.....	32
7.17	Identifier String.....	33
8	Application Flows and Policies	34
8.1	Application Flows.....	34
8.2	Application Flow Specifications and Application Flow Criteria	35
8.3	Application Flow Specification Groups	36
8.4	Policies and Policy Criteria	36
9	SD-WAN Virtual Connection (SWVC) Service Attributes	39
9.1	SWVC Identifier Service Attribute	40
9.2	SWVC List of SWVC End Points Service Attribute.....	41

9.3	SWVC List of UCSs Service Attribute	41
9.4	SWVC Reserved Prefixes Service Attribute	41
9.5	SWVC List of Zones Service Attribute	42
9.6	SWVC List of Virtual Topologies Service Attribute	43
9.6.1	<i>vtType=multipoint-to-multipoint</i>	44
9.6.2	<i>vtType=rooted-multipoint</i>	45
9.7	SWVC Performance Time Intervals Service Attribute	46
9.8	SWVC Service Performance Objectives Reporting Periods Service Attribute	46
9.9	SWVC List of Rate Limiters Service Attribute	46
9.10	SWVC List of Security Policies Service Attribute	47
9.11	SWVC List of Application Flow Business Importance Levels Service Attribute	47
9.12	SWVC List of Policies Service Attribute	48
9.12.1	Policy Criteria specification and interaction	51
9.12.2	Ingress Policy Criteria	52
9.12.3	Egress Policy Criteria	65
9.13	SWVC List of Application Flow Specification Groups Service Attribute	66
9.14	SWVC List of Application Flow Specifications Service Attribute	67
10	SD-WAN Virtual Connection (SWVC) End Point Service Attributes	73
10.1	SWVC End Point Identifier Service Attribute	73
10.2	SWVC End Point Associated UNI Service Attribute	73
10.3	SWVC End Point List of UCS End Points Service Attribute	73
10.4	SWVC End Point Policy Map Service Attribute	74
10.4.1	Ingress Policy Assignment	75
10.4.2	Egress Policy Assignment	77
10.4.3	Examples of Ingress Application Flows and Policy Assignment	78
10.5	SWVC End Point Network Address Translation Service Attribute	79
11	SD-WAN UNI Service Attributes	81
11.1	SD-WAN UNI Identifier Service Attribute	81
11.2	SD-WAN UNI L2 Interface Service Attribute	82
11.3	SD-WAN UNI L2 Maximum Frame Size Service Attribute	83
11.4	SD-WAN UNI IPv4 Connection Addressing Service Attribute	83
11.5	SD-WAN UNI IPv6 Connection Addressing Service Attribute	85
11.6	SD-WAN UNI Routing Protocols Service Attribute	87
11.6.1	Static	89
11.6.2	BGP	89
11.6.3	OSPF	92
11.7	SD-WAN UNI Availability Objective Service Attribute	93
11.8	SD-WAN UNI Mean Time To Repair Objective Service Attribute	94
12	Attributes for Underlay Connectivity Services	98
12.1	SD-WAN UCS Service Attributes	98
12.1.1	UCS Identifier Service Attribute	99
12.1.2	UCS Type Service Attribute	99
12.1.3	UCS Billing Method Service Attribute	99
12.2	SD-WAN UCS UNI Service Attributes	99
12.2.1	UCS UNI Identifier Service Attribute	100
12.3	SD-WAN UCS End Point Service Attributes	100

12.3.1 UCS End Point Identifier Service Attribute..... 101

12.3.2 UCS End Point Backup Service Attribute 101

12.3.3 UCS End Point Breakout Service Attribute..... 101

13 Performance Metrics 103

13.1 Qualified Packets 103

13.2 One-Way Packet Delay 103

13.3 One-Way Mean Packet Delay Performance Metric 104

13.4 One-Way Mean Packet Delay Variation Performance Metric 104

13.5 One-Way Packet Loss Ratio Performance Metric..... 104

14 References 106

Appendix A Processing Application Flows (Informative) 109

A.1 Process View 109

A.2 Ingress Packet Flow View 111

Appendix B SD-WAN Use Cases (Informative) 112

B.1 Hybrid WAN 112

B.2 Dual Internet WAN 113

Appendix C PERFORMANCE Policy Criterion Examples (Informative)..... 114

Appendix D Major changes from MEF 70.1 to MEF 70.2 (Informative) 116

Appendix E Acknowledgements (Informative) 117

List of Figures

Figure 1 – Diagram Conventions	9
Figure 2 – Ingress and Egress	15
Figure 3 – Components of an SD-WAN Service	17
Figure 4 – Multiple SD-WAN UNIs.....	19
Figure 5 – Relationship of Service Components	21
Figure 6 – TVCs.....	22
Figure 7 – Different Virtual Topologies in an SWVC.....	25
Figure 8 – Internet Breakout.....	26
Figure 9 – Example of Unnamed Rate Limiters	29
Figure 10 – Example of Named Rate Limiters	29
Figure 11 – Intended Behavior of a Rate Limiter	30
Figure 12 – Examples of Application Flows	36
Figure 13 – Zones and Policy Assignment Examples	37
Figure 14 – Example of two Multipoint-to-Multipoint Virtual Topologies	44
Figure 15 – Example of Rooted-multipoint Virtual Topology	45
Figure 16 – PERFORMANCE Policy Criterion Parameters	62
Figure 17 – Precedence for Policy Assignment.....	76
Figure 18 – Assigning Ingress Policies.....	78
Figure 19 – BGP Configurations Examples.....	90
Figure 20 – Examples of Outages and Exception Periods.....	96
Figure 21 – Application Flows, Policies, and Forwarding	109
Figure 22 – Ingress Packet Flow.....	111
Figure 23 – Use Case: Hybrid WAN	112
Figure 24 – Use Case: Dual Internet WAN	113

List of Tables

Table 1 – Terminology.....	6
Table 2 – Abbreviations.....	6
Table 3 – Numerical Prefix Conventions.....	8
Table 4 – Summary of SWVC Service Attributes	40
Table 5 – Policy Criteria – Support Required.....	49
Table 6 – Policy Criteria – Support Recommended	50
Table 7 – Performance Metrics	60
Table 8 – Application Flow Criteria – Support Required.....	69
Table 9 – Application Flow Criteria – Support Recommended	71
Table 10 – Summary of SWVC End Point Service Attributes	73
Table 11 – Summary of SD-WAN UNI Service Attributes	81
Table 12 – Summary of SD-WAN UCS Service Attributes.....	98
Table 13 – Summary of SD-WAN UCS UNI Service Attributes.....	100
Table 14 – Summary of SD-WAN UCS End Point Service Attributes	100
Table 15 – Example measured values of Performance Metrics at time tI	114
Table 16 – Example measured values of Performance Metrics at time $tI + \Delta t$	114
Table 17 – Example measured values of Performance Metrics at time $tI + 2\Delta t$	115
Table 18 – Example measured values of Performance Metrics at time $tI + 3\Delta t$	115

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.

- AT&T
- Bell Canada
- Cisco Systems
- Orange Business
- Versa Networks

2 Abstract

The SD-WAN Service Attributes and Service Framework Standard defines the externally visible behavior of a MEF SD-WAN Service. A Service deployment is based on an agreement between an SD-WAN Subscriber (the buyer) and an SD-WAN Service Provider (the seller) that includes agreement on the values of a set of SD-WAN Service Attributes defined in this document.

This document includes:

SD-WAN Key Concepts – description of important components and concepts, i.e., the building blocks that are used to define and describe a MEF SD-WAN Service.

SD-WAN Service Attributes – the enumeration and description of the information that is agreed to between the SD-WAN Subscriber and the SD-WAN Service Provider. The values of these Service Attributes are determined by agreement between the Subscriber and Service Provider, subject to constraints imposed by this standard.

SD-WAN Service Framework – A framework for defining instances of an SD-WAN Service based on the definitions, service elements, and Service Attributes included in the document.

This document supersedes and replaces MEF 70.1, *SD-WAN Service Attributes and Service Framework* [36].

3 Terminology and Abbreviations

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

In addition, terms defined in MEF 61.1 [32] are included in this document by reference and are not repeated in the table below. Terms marked with * are adapted from terms in MEF 10.4 [30] or MEF 61.1 [32].

Term	Definition	Reference
Application Flow	A sequence of IP Packets that ingress at an SD-WAN UNI or are directed towards an SD-WAN UNI from a UCS that: <ol style="list-style-type: none"> 1. Have the same direction (ingress or egress), and 2. Match the same Application Flow Specification, and 3. Are in the same Zone (i.e., have source IP Addresses in the same Zone or are all in the Zone <i>Internet</i>) 	This document
Application Flow Criterion	A specific condition for matching an IP Packet such as a field/value pair or identification by an algorithm or heuristic.	This document
Application Flow Specification Group	A named set of Application Flow Specifications.	This document
Application Flow Specification	A named set of Application Flow Criteria.	This document
Bandwidth Flow	A set of Application Flows that are subject to a particular set of bandwidth constraints.	This document
Egress Application Flow	An Application Flow consisting of IP Packets directed toward an SD-WAN UNI by the Service Provider.	This document
Egress IP Packet	An IP Packet transmitted to the Subscriber Network by the Service Provider at an SD-WAN UNI.	This document *
Egress Policy	A Policy that is assigned to Egress Application Flows.	This document
Egress Policy Criterion	A Policy Criterion that is applicable to Egress Application Flows.	This document
Egress UNI	For a given IP Packet flowing over the SD-WAN Service, the SD-WAN UNI where the IP Packet is received by the Subscriber Network from the Service Provider Network.	This document
Identifier String	A string composed of a well-defined set of octets that is used to name various components of the SD-WAN Service for management purposes.	This document
Ingress Application Flow	An Application Flow consisting of Ingress IP Packets at an SD-WAN UNI.	This document
Ingress IP Packet	An IP Packet received from the Subscriber Network by the Service Provider at an SD-WAN UNI.	This document *
Ingress Policy	A Policy that is assigned to Ingress Application Flows.	This document

Term	Definition	Reference
Ingress Policy Criterion	A Policy Criterion that is applicable to Ingress Application Flows.	This document
Ingress UNI	For a given IP Packet flowing over the SD-WAN Service, the SD-WAN UNI where the IP Packet is received by the Service Provider Network from the Subscriber Network.	This document
Internet Access Service	Public Internet connectivity service purchased by an IP Subscriber from an Internet Service Provider.	This document
Internet Breakout	The forwarding of IP Packets in Application Flows, based on Policy, to Internet destinations via Internet Access UCSs.	This document
Local Internet Breakout	Internet Breakout in which Ingress IP Packets are forwarded over Internet Access UCSs connected to the SD-WAN Edge where the Ingress SD-WAN UNI is located.	This document
Path	A single TVC or a sequence of TVCs that connect a pair of SD-WAN Edges over which IP Packets are forwarded between UNIs located at the SD-WAN Edges.	This document
Performance Metric	One of several performance-related properties of a Path that can be measured.	This document
Policy	A named list of Policy Criteria that can be assigned to an Application Flow and that determines how an SD-WAN Service handles IP Packets in the Application Flow.	This document
Policy Criterion	A criterion that describes a specific objective or constraint on IP Packet handling.	This document
Qualified Packet	An Ingress IP Packet that complies with specific criteria and for which Performance Metrics are defined.	This document
Rate Limiter	An abstract component of a network that measures the information rate of a Bandwidth Flow and takes actions to conform with the specified bandwidth constraints.	This document
SD-WAN Edge	A set of network functions (physical or virtual) that are located between the SD-WAN UNI(s) and the Underlay Connectivity Service UNI(s).	This document
SD-WAN Service	An overlay connectivity service that optimizes transport of IP Packets over one or more other connectivity services (Underlay Connectivity Services) by recognizing applications (Application Flows) and determining forwarding behavior (by applying Policies). MEF SD-WAN Services are specified using Service Attributes defined in this MEF Standard.	This document
SD-WAN Service Provider	A Service Provider for an SD-WAN Service	This document
SD-WAN Subscriber	A Subscriber of an SD-WAN Service	This document
SD-WAN UNI	SD-WAN User Network Interface	This document
SD-WAN User Network Interface	The demarcation point between the responsibility of the SD-WAN Service Provider and the SD-WAN Subscriber.	This document

Term	Definition	Reference
SD-WAN Virtual Connection	An association of SD-WAN Virtual Connection End Points in an SD-WAN Service that provides the logical construct of a L3 Virtual Private Routed Network for a Subscriber.	This document
SD-WAN Virtual Connection End Point	A logical construct associated with an SD-WAN UNI where Policies are applied to Ingress and Egress Application Flows.	This document
Security Policy	A set of parameters that are agreed between the Subscriber and Service Provider (as part of the SWVC List of Policies Service Attribute) and that specify which Security Functions are to be applied to an Application Flow.	MEF 88 [38]
Service Performance Objective	The SD-WAN UNI Availability Objective or SD-WAN UNI Mean Time To Repair Objective	This document
Service Provider	An organization that provides services to Subscribers. In this document, "Service Provider" means "SD-WAN Service Provider".	This document *
Service Provider Network	As seen from the point-of-view of the Subscriber, the aggregation of Underlay Connectivity Services (which may be provided by different organizations), TVCs, SD-WAN Edges, controllers, and orchestrators used to deliver an SD-WAN Service to a Subscriber. In this document, "Service Provider Network" means "SD-WAN Service Provider Network".	This document
Subscriber	The end-user of a service. In this document, "Subscriber" should be read as meaning "SD-WAN Subscriber".	This document *
Subscriber Location	Any place where there is an SD-WAN UNI.	This document
Subscriber Network	A network belonging to a given Subscriber (or other organization authorized by the Subscriber) that is connected to the Service Provider Network at one or more SD-WAN UNIs.	This document *
SWVC End Point	SD-WAN Virtual Connection End Point	This document
Tunnel Virtual Connection	A point-to-point forwarding relationship between two SD-WAN Edges that associates: <ul style="list-style-type: none"> two UCS End Points in a given single Underlay Connectivity Service that is not an Internet Access Service, or a UCS End Point in each of two Internet Access Underlay Connectivity Services and has a well-defined set of packet delivery characteristics (e.g., delay, security, bandwidth, etc.)	This document
Underlay Connectivity Service	A service providing connectivity between two or more Subscriber Locations, or between a Subscriber Location and the Internet, over which an SD-WAN Service is provided; for example, a private IP Service or a Carrier Ethernet service.	This document
Underlay Connectivity Service End Point	A logical construct at the Underlay Connectivity Service UNI to which a distinct subset of IP Packets passing over the Underlay Connectivity Service UNI is mapped.	This document

Term	Definition	Reference
Underlay Connectivity Service Provider	An organization that provides an Underlay Connectivity Service to a Subscriber or SD-WAN Service Provider.	This document
Underlay Connectivity Service UNI	A UNI in an Underlay Connectivity Service	This document
Virtual Topology	A specific definition of the allowed forwarding behavior between a defined set of SD-WAN UNIs in an SWVC that can be applied to certain Application Flows by Policy.	This document
Zone	A subset of the IP hosts in the Subscriber Network identified by a set of IP Prefixes.	This document
Zone-wide Policy	A Policy that is applied to IP Packets in a Zone for specified Application Flows at all UNIs unless overridden by another Policy.	This document

Table 1 – Terminology

Abbreviation	Definition	Reference
SD-WAN	Software Defined Wide Area Network (see SD-WAN Service)	This document
SWVC	SD-WAN Virtual Connection	This document
TVC	Tunnel Virtual Connection	This document
UCS / UCSs	Underlay Connectivity Service / Services	This document
UCS UNI	Underlay Connectivity Service User Network Interface	This document
UNI	Short for User Network Interface. In this document, UNI without a modifier (such as “UCS UNI”) means SD-WAN UNI.	This document

Table 2 – Abbreviations

4 Compliance Levels

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in BCP 14 (RFC 2119 [7], RFC 8174 [27]) 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.

5 Document Conventions

This section describes the numerical, notational, and diagrammatic conventions used in the document.

5.1 Numerical Prefix Conventions

This Standard uses the prefix notation to indicate multiplier values as shown 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 Prefix Conventions

5.2 Notational Conventions

There are several places in the document where mathematical, set, and other notational structures are specified. These items are specified using different types of brackets as follows:

- $\langle xxx, yyy \rangle$ – wide angle brackets are used to surround n-tuples
- $[xxx, yyy, zzz]$ – square brackets are used to surround lists
- $\{ a, b, c \}$ – braces are used to surround sets

These brackets can be composed in many ways. For example:

$[\langle x1, y1 \rangle, \langle x2, y2 \rangle, \langle x3, y3 \rangle]$ = a list of 2-tuples

$\langle x1, [\langle a1, b1, c1 \rangle, \langle a2, b2, c2 \rangle, \langle a3, b3, c3 \rangle] \rangle$ = a 2-tuple where the second element is a list of 3-tuples

5.3 Diagram Conventions

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




Figure 1 – Diagram Conventions

- SD-WAN Edges are represented by rectangles with rounded corners. The edges are thicker gray lines, and they are white or light gray inside.
- SD-WAN User Network Interfaces (UNIs) are solid red rectangles. They are usually shown in conjunction with an SD-WAN Edge.
- SWVC End Points are represented as circles (sometimes as ovals in order to fit text) and are usually shown in conjunction with the SD-WAN Edge and SD-WAN UNI.
- The UCS UNI is shown as a solid blue rectangle and is usually shown on the opposite side of the SD-WAN edge from the SD-WAN UNI.
- The UCS End Point is a small circle attached to the UCS side of the UCS UNI.
- A TVC is represented as a line of any color with diamonds at both ends.

6 Introduction

An SD-WAN Service is a connectivity service that creates an overlay network over one or more Underlay Connectivity Services. An SD-WAN Service enables orchestrated connectivity that recognizes Application Flows at ingress to the Service and forwards them between SD-WAN User Network Interfaces based on Policies that are applied to them. An SD-WAN Service provides the logical construct of a Virtual Private Routed Network (see RFC 2764 [13]) for the Subscriber that transports IP Packets between Subscriber Locations. Additionally, an SD-WAN Service may provide connectivity between SD-WAN UNIs and the public Internet.

Since the SD-WAN Service can use multiple and possibly disparate Underlay Connectivity Services, it can offer more differentiated service delivery capabilities than a single Ethernet or IP Service.

It is important to distinguish between the SD-WAN Service and the Underlay Connectivity Services over which the SD-WAN Service operates. An SD-WAN Service Provider is responsible for providing the SD-WAN Service to the Subscriber. The SD-WAN Service Provider may be responsible for none, some, or all the Underlay Connectivity Services. The SD-WAN Service may use Underlay Connectivity Services that the Subscriber has obtained from UCS Service Providers other than the SD-WAN Service Provider.

An SD-WAN service is aware of, and forwards traffic based on, Application Flows and Policies (see section 7.11). The Service agreement includes specification of Application Flows—IP Packets that match a set of criteria—and Policies that describe requirements and constraints on the forwarding of the Application Flows.

A benefit of an SD-WAN Service is the ability to adjust aspects of the service in near real time to meet business needs. These business needs are expressed in SD-WAN Policies as defined in this document. The Service Provider monitors the performance of the available forwarding options and modifies how packets in each Application Flow are forwarded to achieve the desired behavior.

This document defines Service Attributes that describe the externally visible behavior and operation of an SD-WAN Service as experienced by the Subscriber. It also defines Service Attributes for the Underlay Connectivity Services (and related components), which are necessary to define the behavior of SD-WAN Policies.

Several SD-WAN use cases are described in Appendix B.

6.1 Document Scope

MEF 70.2 includes definition and description of:¹

- SD-WAN Service components and SD-WAN Service framework
- SD-WAN Service functionality visible to the Subscriber
- Service Attributes for SD-WAN Virtual Connection (SWVC), SWVC End Point, and SD-WAN UNI
- Policies and the Policy Criteria that compose them

¹ A list of major differences between MEF 70.1 and MEF 70.2 is provided in Appendix C.

- Definition of a specific set of Policy Criteria
- Application Flows, Application Flow Specifications, and the Application Flow Criteria that compose the Application Flow Specifications
- The SD-WAN UNI and details of Subscriber connection to the SD-WAN Service
- Service Attributes to describe key characteristics of an Underlay Connectivity Services (UCS), Underlay Connectivity Service UNIs, and Underlay Connectivity Service End Points from an SD-WAN Service perspective
- The assignment of Policies to Ingress and Egress Application Flows at an SWVC End Point
- Virtual Topologies that can be assigned to Ingress Application Flows by Policy
- Partitioning the Subscriber's IP hosts into Zones
- Key characteristics of Tunnel Virtual Connections

6.2 Characteristics of an SD-WAN Service

A MEF SD-WAN Service has the following characteristics:

- The Subscriber connects to the SD-WAN Service at an SD-WAN UNI.
- The basic unit of transport at the SD-WAN UNI is an IP Packet.
- The SD-WAN Service provides a layer 3, IP routed network.
- Ingress IP Packets at the UNI are classified, based on the IP Packet contents, into Application Flows.
- The SD-WAN Service utilizes one or more Underlay Connectivity Services.
- Paths across the Underlay Connectivity Services are selected dynamically for Application Flows to conform with assigned Policies.
- Policies can specify performance goals for each Application Flow.
- SWVC topologies are defined by Policies and IP forwarding constraints.
- An SD-WAN Service can offer encryption between SD-WAN Edges.
- Forwarding of an Application Flow can be blocked at an SWVC End Point by Policy.
- Each Application Flow can, by Policy, be subject to a bandwidth commitment and limit. Policies can be defined to allow sharing of bandwidth commitments and limits between Application Flows.
- An SD-WAN Service typically provides a Subscriber web portal and/or API that exposes network health, performance, and application information. The portal/API may also allow the Subscriber to modify aspects of the SD-WAN service such as defining Application Flow Specifications and creating/modifying Policies.
- An SD-WAN Service aligns with the concepts of MEF LSO principles including Service Orchestration.

6.3 Organization of the Standard

The remainder of the document is organized as follows:

- Definitions and key concepts are detailed in section 7.
- An overview of Application Flows and Policies is provided in section 8.
- Service Attributes for the SD-WAN Virtual Connection (SWVC) are described in section 9.

- Service Attributes for the SD-WAN Virtual Connection End Point are described in section 10.
- Service Attributes for the SD-WAN UNI are described in section 11.
- Service Attributes for Underlay Connectivity Service, UCS UNI, and UCS End Point are described in sections 12.1, 12.2, and 12.3 respectively.
- Performance Metrics used in the PERFORMANCE Policy Criterion are defined in section 13.

7 Key Concepts and Definitions

This section provides definitions, key concepts, and overviews of the components of a MEF SD-WAN Service. An attempt has been made to present the concepts in an order that minimizes forward references, i.e., building from the most basic concepts.

Section 7 includes the following subsections:

- Service Attributes (section 7.1)
- SD-WAN Subscriber and SD-WAN Service Provider (section 7.2)
- Subscriber Network and Service Provider Network (section 7.3)
- SD-WAN UNI (section 7.4)
- SWVC End Point (section 7.5)
- SWVC and SD-WAN Service Components Overview (section 7.6)
- Underlay Connectivity Service, UCS UNI, and UCS End Point (section 7.7)
- SD-WAN Edge (section 7.8)
- SD-WAN Service Framework (section 7.9)
- Tunnel Virtual Connection (section 7.10)
- Application Flow, Application Flow Specification, and Policy (section 7.11)
- Zone (section 7.12)
- Virtual Topology (section 7.13)
- Internet Breakout (section 7.14)
- Rate Limiter and Bandwidth Flow (section 7.15)
- SD-WAN IP Packet Delivery (section 7.16)
- Identifier String (section 7.17)

7.1 Service Attributes

The behavior of MEF Services, such as SD-WAN, is specified using Service Attributes. A Service Attribute captures specific information agreed on between the Service Provider and the Subscriber of a MEF Service, and it describes some aspect of the service behavior. How such an agreement is reached, and the specific values agreed upon, might have an impact on the price of the service or on other business or commercial aspects of the relationship between the Subscriber and the Service Provider; these details are outside the scope of this document. Some examples of how an agreement could be reached are given below, but this is not an exhaustive list.

- The Service Provider mandates a particular value.
- The Subscriber selects from a set of options specified by the Service Provider.
- The Subscriber requests a particular value, and the Service Provider accepts it.
- The Subscriber and the Service Provider negotiate to reach a mutually acceptable value.

Service Attributes describe the externally visible behavior of the service as experienced by the Subscriber, including the definitions of major SD-WAN capabilities such as Application Flow Specifications, Policies, Zones, Virtual Topologies, etc., that affect how traffic is handled within the SD-WAN Service. However, they do not constrain how the Service Provider implements the service, nor how the Subscriber implements their network. The Subscriber and the Service Provider agree upon the initial value for each Service Attribute in advance of the service

deployment. The Subscriber and the Service Provider may subsequently agree on changes to the values of certain Service Attributes. This document does not constrain how such agreement is reached; for example, if the Service Provider allows the Subscriber to select an initial value from a pre-determined set of values, they might further allow them to change their selection at any time during the lifetime of the service.

7.2 SD-WAN Subscriber and SD-WAN Service Provider

This document deals, primarily, with two organizations—the SD-WAN Subscriber and the SD-WAN Service Provider. The SD-WAN Subscriber is the organization that uses services described by the Service Attributes specified in this document. The SD-WAN Service Provider is the organization that provides those services.

There is no restriction on the type of organization that can act as a Subscriber. For example, a Subscriber can be an enterprise, a mobile operator, an IT system integrator, or a government department. At its most basic, an SD-WAN Service provides connectivity for IP Packets between different parts of the Subscriber Network or between the Subscriber Network and the Public Internet using Internet Breakout (described in section 7.14).

The remainder of this document uses “Service Provider” to refer to the SD-WAN Service Provider and “Subscriber” to refer to the SD-WAN Subscriber.

7.3 Subscriber Network and Service Provider Network

The Subscriber Network is defined as the network belonging to the Subscriber (or another organization authorized by the Subscriber) that is connected to the Service Provider Network at UNIs.

The Service Provider Network represents the set of components and logical constructs used by the Service Provider to provide an SD-WAN Service. It includes SD-WAN Edges (section 7.8), Underlay Connectivity Services (section 7.7), Tunnel Virtual Connections (section 7.10), and management and control functions.

The term “Service Provider Network” indicates that the SD-WAN Service Provider integrates all of these components into the SD-WAN Service offering but does not make a statement about the owner or operator of any given component—all components may be owned and operated by the Service Provider, or some components (or even all) may be owned and operated by one or more other organizations.

The Service Provider Network can be completely opaque, that is, the Subscriber connects to the Service Provider Network at the UNIs, and the SD-WAN Service provides the desired connectivity, but the Subscriber has no insight into any of the underlying components. Alternatively, the Subscriber may have visibility into certain components in the Service Provider Network; for example, the Subscriber can agree with the Service Provider that some of the Subscriber’s existing (i.e., previously contracted) Underlay Connectivity Services (see section 7.7) are to be used with the SD-WAN Service. In that case, these Underlay Connectivity Services are known to the Subscriber; however, other underlying components of the SD-WAN Service may remain opaque.

7.4 SD-WAN UNI

An SD-WAN User Network Interface, or SD-WAN UNI, is the demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber. The SD-WAN UNI is on the boundary between the Subscriber Network and the Service Provider Network (see section 7.3). The basic unit of transport at the SD-WAN UNI is an IP Packet.

In the remainder of this document when the abbreviation “UNI” is used without a modifier, it refers to the SD-WAN UNI (references to the UCS UNI always include “UCS”).

Formally, the UNI is a demarcation point. The term UNI is often used to also refer to the network connection between the Subscriber Network and the Service Provider Network, but the actual location of the reference point is important because it defines where the Service Provider’s responsibility starts and also because service performance is defined from UNI to UNI.

[R1] An SD-WAN UNI **MUST** be dedicated to a single Subscriber.

[R2] An SD-WAN UNI **MUST** be dedicated to a single Service Provider.

An IP Packet that crosses the UNI from the Subscriber to the Service Provider is called an *Ingress* IP Packet, and the UNI is the *Ingress UNI* for that IP Packet. Similarly, an IP Packet that crosses the UNI from Service Provider to the Subscriber is called an *Egress* IP Packet, and the UNI is the *Egress UNI* for that IP Packet. These are shown in the following diagram:

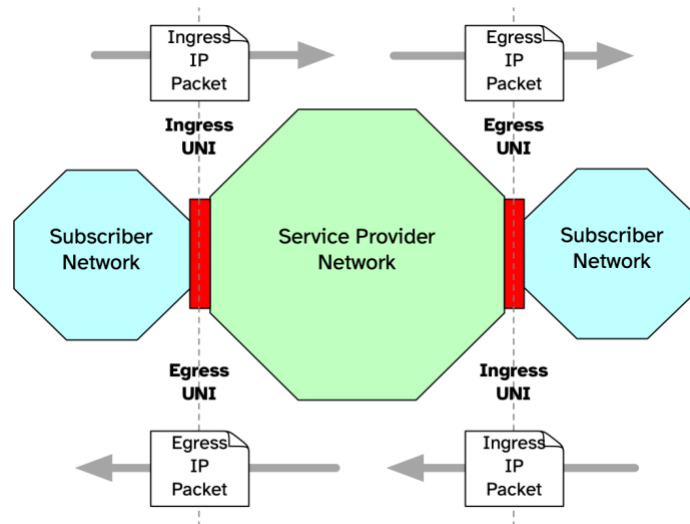


Figure 2 – Ingress and Egress

Terminology note: Since SD-WAN is an overlay network service, the terms *ingress* and *egress* can also be (and often are) applied to IP packets crossing the Underlay Connectivity Service UNIs. This can create ambiguity and confusion. In this document their use *only* refers to packets crossing the SD-WAN UNI as shown in Figure 2.

An SD-WAN UNI connects to one or more IPv4 or IPv6 Subscriber subnets (see sections 11.4 and 11.5).

7.5 SWVC End Point

An SWVC End Point is a logical construct associated with each UNI (one SWVC End Point per UNI) that encapsulates the Service Attributes that assign Ingress Policies to Ingress Application Flows and Egress Policies to Egress Application Flows.

[R3] An SD-WAN UNI **MUST** be associated with, at most, one SWVC End Point.

Requirement [R3] allows a UNI without an SWVC End Point, but no SD-WAN traffic can be forwarded at this UNI. This may be a transitory state during SD-WAN Service creation or modification.

7.6 SWVC and SD-WAN Service Components Overview

An SD-WAN Service comprises several logical constructs enumerated in this section. The first set of constructs are specified by Service Attributes that are agreed to between the Service Provider and the Subscriber. The second set of constructs are used to describe various aspects of the Service capabilities and behavior but are not specified by Service Attributes.

MEF SD-WAN Services are specified by Service Attributes for the following:

- SD-WAN Virtual Connection (SWVC) (this section and section 9)
- SD-WAN Virtual Connection End Point (sections 7.5 and 10)
- SD-WAN UNI (sections 7.4 and 11)
- Certain aspects of the Underlay Connectivity Services (UCS), UCS UNIs, and UCS End Points (sections 7.7 and 12)

The SD-WAN Service constructs that are not described by Service Attributes and may (or may not) be visible to the Subscriber are listed below:

- Subscriber Network (clearly this is visible to the Subscriber) (section 7.3)
- Tunnel Virtual Connection (TVC) (section 7.10)
- SD-WAN Edge (section 7.8)

The SWVC is an association of SWVC End Points located at SD-WAN UNIs. The SWVC defines the logical connectivity of the SD-WAN Service as viewed by the Subscriber. The SWVC accepts IP Packets at Ingress UNIs and delivers them to Egress UNIs (or to the Internet in the case of Internet Breakout).

Figure 3 shows the SD-WAN Service constructs described above.

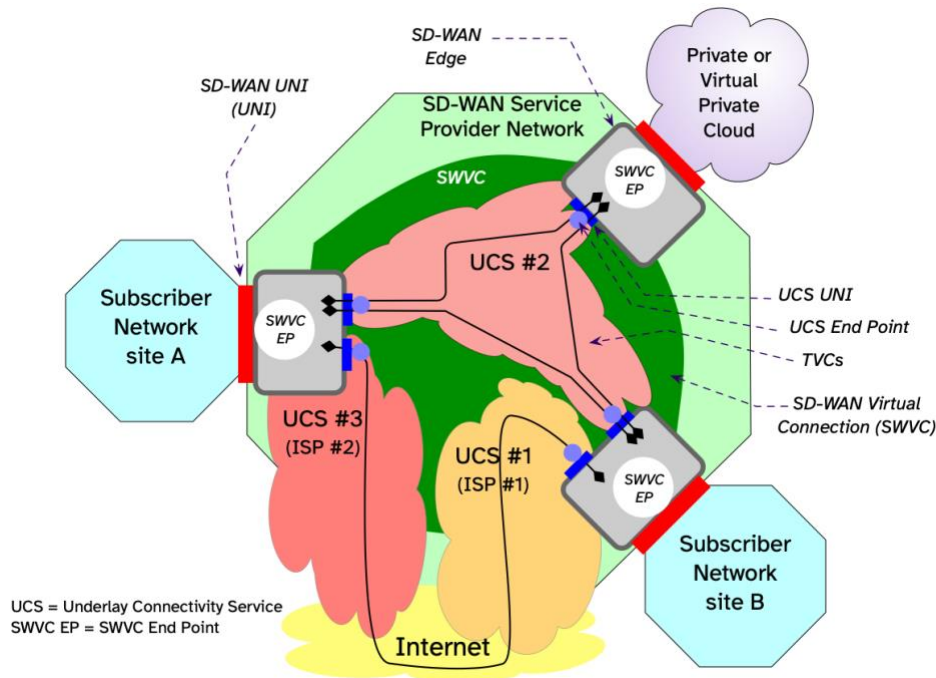


Figure 3 – Components of an SD-WAN Service

Note that in Figure 3, one of the Subscriber sites is a connection to a Private or Virtual Private Cloud, which may not be located at the Subscriber’s physical location. However, the SD-WAN Edge and its components (SD-WAN UNI, SWVC End Point, etc.) at this Subscriber Location operate no differently than at any of the other Subscriber Locations.

Each Underlay Connectivity Service terminates at a service demarcation point, i.e., a UCS UNI (since they are also services), which is shown in the diagram. Depending on the type of Underlay Connectivity Service, this could be an Ethernet UNI (as defined in MEF 10.4 [30], an IP UNI (as defined in MEF 61.1 [32], including Internet Access as defined in MEF 69.1 [35]), an L1 UNI (as defined in MEF 63 [34]), or analogous service demarcation for other non-MEF services. Since the UCS UNI can provide access to multiple services (consider a UCS UNI with multiple Carrier Ethernet services), the UCS End Point is the logical construct that associates a specific UCS with the UCS UNI (see section 7.7). Note that Figure 3 shows a case where each UCS UNI provides access to a single UCS, and therefore has only a single UCS End Point.

7.7 Underlay Connectivity Service, UCS UNI, and UCS End Point

An SD-WAN Service is an overlay network service that operates over one or more Underlay Connectivity Services. Underlay Connectivity Services are network services that provide connectivity between Subscriber Locations or between a Subscriber Location and the Internet.

Underlay Connectivity Services can include a variety of services, including (but not limited to) Ethernet Services (as defined in MEF 6.3 [29]), IP Services (defined in MEF 61.1 [32]), Internet Access Services (defined in MEF 69.1 [35]), and L1 Connectivity Services (defined in MEF 63 [34]). Access to these Underlay Connectivity Services can be via a variety of networking technologies, such as DSL, HFC, LTE, fiber, Wi-Fi, Ethernet; and the transport can be based on Ethernet switching, IP routing, MPLS, or other technologies.

Underlay Connectivity Services have several characteristics that can be used by Policies that determine the forwarding of Application Flows within the SD-WAN Service.

- An Underlay Connectivity Service is referred to as *Private* or *Public*. A Public UCS is an Internet Access Service such as described in MEF 69.1 [35]. (See section 12.1.2.) A Private UCS is any UCS that is not a Public UCS.
- Cost for usage of an Underlay Connectivity Service is *flat-rate* or *usage-based*. Flat rate means that a given amount of service bandwidth is billed at a fixed amount for the billing period, e.g., \$50/month for 100 Mb/s. Usage-based means that service is billed based on the amount of data that is transmitted or received, e.g., £10/GB. More complex charging structures are also possible. (See section 12.1.3.)
- The Service Provider and Subscriber can agree that an Underlay Connectivity Service is designated as a *Backup* UCS at an SD-WAN Edge (UCS UNI). See section 9.12.2.7 for more detail about the use of a UCS designated as *Backup*. (Also, see section 12.3.2.)
- As with all communications services, Underlay Connectivity Services have bandwidth limitations and performance characteristics that can affect services that run over them.

A UCS is a service with a UCS Subscriber and a UCS Service Provider, and it is likely that there are service parameters and attributes that are agreed on by those parties (examples might be UCS Bandwidth and IP Addressing), but that UCS service agreement is separate from the SD-WAN Service. See section 12 for additional details on UCS Service Attributes.

SD-WAN Services are frequently deployed over multiple, and often disparate, Underlay Connectivity Services. Multiple Underlay Connectivity Services with different performance and cost characteristics (e.g., an IP-VPN over MPLS Network versus an Internet Access Service) can be used to provide cost benefits, resiliency, and differentiated transport.

Underlay Connectivity Services can be provided by the SD-WAN Service Provider on its own network or over the networks of other network operators (including the Internet). Underlay Connectivity Services arranged independently by the Subscriber can also be used by the SD-WAN Service. Some details of the Underlay Connectivity Service are agreed to by the Subscriber and Service Provider as values of UCS, UCS UNI, and UCS End Point Service Attributes specified in this document. The Subscriber also communicates other details of the UCS to the SD-WAN Service Provider in this case as described in section 12. The division of responsibility for UCS services between the Subscriber and the SD-WAN Service Provider is outside the scope of this document.

The UCS has a demarcation point, a UCS UNI, that delineates the responsibility of the UCS Subscriber and the UCS Service Provider.

Many MEF Services, such as Carrier Ethernet Services, support *service multiplexing* at their corresponding UNIs, i.e., a single UNI providing access to multiple services and therefore multiple service end points. For example, a Carrier Ethernet UNI at the headquarters could provide access to an EVP-LAN service connecting to all engineering facilities, and to an EVPL service from the headquarters to the datacenter. To accommodate this, the UCS UNI can likewise support service multiplexing, i.e., multiple UCS End Points at the UCS UNI, and each UCS End Point selects packets from one UCS at the UCS UNI and directs packets towards the one particular UCS at the

UCS UNI. Also, the UCS attributes that can vary at each UCS UNI are associated with a UCS End Point. In the example above, the headquarters location would have two UCS End Points at the UCS UNI, and all the other facilities would have one. There is no requirement that an SD-WAN Service include all the UCS End Points at a UCS UNI.

The UCS End Point is important in the SD-WAN Service context because the characteristics of the UCS, as seen by the SD-WAN Service, are not (necessarily) the same at all UCS UNIs in the UCS. For example, the amount of delivered bandwidth can be different at different UCS End Points.

As noted in section 7.4, the term “UNI” by itself always refers to the SD-WAN UNI. References to the UCS UNI are always “UCS UNI”.

7.8 SD-WAN Edge

The SD-WAN Edge is a set of network functions (physical or virtual) located at the Service Provider side of the SD-WAN UNI, between the SD-WAN UNI(s) and the UCS UNI(s). It is part of the Service Provider Network, but it is commonly located at a Subscriber premises. It is situated between the SD-WAN UNIs, on its Subscriber side, and UCS UNIs of one or more Underlay Connectivity Services on its network side.

The SD-WAN Edge implements functionality that receives ingress IP Packets over the SD-WAN UNIs; determines how they should be handled according to routing information, applicable Policies, other Service Attributes, and information about the UCSs; and if appropriate, forwards them over one of the available UCS UNIs. Similarly, it receives packets over the UCS UNIs and determines how to handle them, including forwarding them over one of the SD-WAN UNIs to the Subscriber Network, if appropriate. The SD-WAN Edge thus implements all the functionality of the SD-WAN service that is not provided by a UCS. This includes Application Flow determination, Policy assignment, and routing/forwarding.

Commonly, each Subscriber Location has an SD-WAN Edge with a single UNI, but an SD-WAN Edge could support multiple UNIs. These UNIs can be separate physical Ethernet connections or a single Ethernet connection with multiple VLANs (see section 11.2) as shown in Figure 4 (or a combination of these approaches). The two examples in Figure 4 are equivalent from the SD-WAN Edge point of view.




Figure 4 – Multiple SD-WAN UNIs

There are several reasons for having multiple physical or logical Ethernet interfaces and their associated UNIs. One reason might just be convenience, for example, connecting the first floor Ethernet and the second floor Ethernet to different ports on the SD-WAN Edge. Another might be

for segregation of traffic, e.g., Ethernet interface #1 might connect to the Guest-WiFi Access Point and Ethernet interface #2 to the corporate router.²

One benefit of supporting multiple UNIs for the same SWVC at an SD-WAN Edge is that Policy assignment is done based on the Ingress UNI³. Therefore, an Application Flow that results from matching an Application Flow Specification and Zone at one Ingress UNI can be assigned a different Policy than an Application Flow that results from matching the same Application Flow Specification and Zone at another Ingress UNI.

7.9 SD-WAN Services Framework

An SD-WAN Service is a connectivity service, offered by a Service Provider, that optimizes transport of IP Packets over one or more Underlay Connectivity Services by recognizing Application Flows at ingress to the Service and forwarding them based on Policies that are applied to them.

A MEF SD-WAN Service consists of:

- Exactly one SWVC with a corresponding set of SWVC Service Attributes (see section 9).
- One⁴ or more UNIs where the Subscriber accesses the SD-WAN Service, each with a corresponding set of UNI Service Attributes (see section 11).
- Exactly one SWVC End Point for the SWVC at each of those UNIs, where each SWVC End Point has a corresponding set of SWVC End Point Service Attributes (see section 10).
- One or more Underlay Connectivity Services (section 12.1) and corresponding UCS UNIs (section 12.2) and UCS End Points (section 12.3) and corresponding Service Attributes for each of them.

There is a one-to-one relationship between an SD-WAN Service and an SWVC. Note that the SWVC, the SWVC End Points, the UNIs, the Underlay Connectivity Services, and the UCS End Points (and their Service Attributes) are specific to a given SD-WAN Service instance.

[R4] Each UCS End Point **MUST** be associated with at most one SWVC.

[R5] Each Private UCS **MUST** be associated with at most one SWVC.

The reason that [R4] and [R5] do not include the UCS UNI is that one UCS UNI could support multiple UCSs (each with a UCS End Point at the UCS UNI) and therefore a UCS UNI could support multiple SWVCs. Note, however, that supporting multiple SWVCs in an SD-WAN Edge is out of scope for this document. In addition, a Public UCS (Internet Access Service) can support multiple SWVCs which is why [R5] is restricted to Private UCSs.

² It is also possible that the UNIs are associated with different SD-WAN Services (SWVCs), but this is not in scope for this document.

³ Policy is assigned at the SWVC End Point, but there is a one-to-one relationship between the UNI and the SWVC End Point.

⁴ An SD-WAN Service with only one UNI is unusual but there are scenarios where it might be appropriate.

The relationships between the various components of the SD-WAN Service and their Service Attributes are shown in the following diagram.

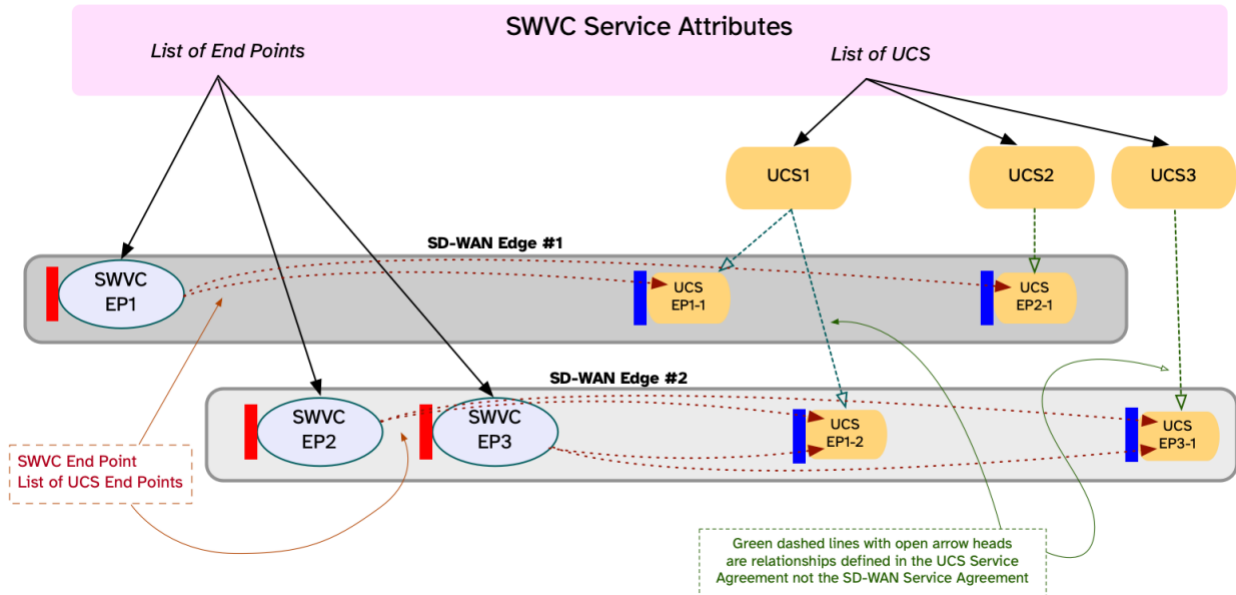


Figure 5 – Relationship of Service Components

An SWVC is an association of SWVC End Points, and the SWVC includes a Service Attribute that identifies the SWVC End Points associated by the SWVC. Each SWVC End Point has a Service Attribute to identify the UNI at which it is located.

Each SWVC also has a list of Underlay Connectivity Services that are used to provide transport for the SWVC. Each UCS has a UCS End Point associated with each SD-WAN Edge that it services. Each UCS End Point associates its UCS with a UCS UNI.

In Figure 5, the SWVC has three SWVC End Points. EP1 is in SD-WAN Edge #1, and EP2 and EP3 are in SD-WAN Edge #2 (each at a different SD-WAN UNI). There are three Underlay Connectivity Services. UCS1 has two UCS End Points, each associated with one of the SD-WAN Edges. UCS2 and UCS3 each has one UCS End Point—they are Internet Access Services. All SWVC End Points in an SD-WAN Edge have access to all UCS End Points in that SD-WAN Edge. Although not shown in the diagram, each UCS UNI could have multiple UCS End Points associated with it. For example, UCS EP1-1 and UCS EP2-1 could have shared the same UCS UNI in SD-WAN Edge #1.

7.10 Tunnel Virtual Connection (TVC)

A Tunnel Virtual Connection (TVC) is a point-to-point forwarding relationship between two SD-WAN Edges that associates either:

- two UCS End Points in a single Underlay Connectivity Service that is not an Internet Access Service, or
- a UCS End Point in each of two Internet Access Underlay Connectivity Services (such as defined in MEF 69.1 [35])

A Path between two SD-WAN UNIs is a single TVC between the SD-WAN Edges where the UNIs are located or a sequence of TVCs that link those SD-WAN Edges.

A TVC has performance and privacy/security characteristics, e.g., delay, bandwidth, encryption, etc.

- [R6] For an IP Packet in an Application Flow to be forwarded directly from one SD-WAN Edge to another SD-WAN Edge, there **MUST** be at least one TVC between the two SD-WAN Edges.

Although the name Tunnel Virtual Connection includes the word “tunnel”, this does not imply a particular implementation or instantiation, but rather that the forwarding relationship has a well-defined set of end points that source and sink packets.

As shown in Figure 6, two SD-WAN Edges can be connected by one TVC or multiple TVCs. Multiple TVCs can be provided between the same two SD-WAN Edges for several reasons, including:

- They traverse different UCSs and can provide resilience
- They traverse different UCSs and can provide additional bandwidth or load balancing
- They have different performance characteristics (e.g., assigned to different UCS classes of service)
- They have different privacy/security characteristics

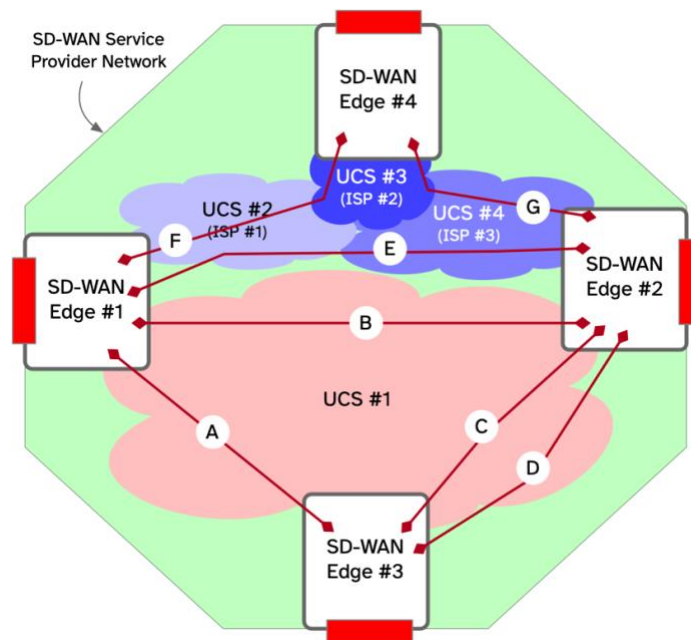


Figure 6 – TVCs

As shown in Figure 6 and based on [R6]:

- TVCs E, F, and G provide connectivity between (respectively):
 - SD-WAN Edges 1–2 (across UCS #2 and UCS #4),
 - SD-WAN Edges 1–4 (across UCS #2 and UCS #3), and
 - SD-WAN Edges 2–4 (across UCS #3 and UCS #4).

- TVCs A and C provide connectivity across UCS #1 between SD-WAN Edges 1 and 3 and 2 and 3.
- TVC B provides connectivity between SD-WAN Edges 1 and 2, which is also provided by TVC E. These two TVCs traverse different UCSs, which means they can provide different characteristics such as more bandwidth or more resiliency or different cost, etc.
- TVC D is between SD-WAN Edge 2 and SD-WAN Edge 3 parallel to C. Both are across the same UCS so there is no additional resiliency, but they could have different classes of service, or one of them might be encrypted, or they could be allocated to support different Application Flows.
- Note that all Paths between the UNI at SD-WAN Edge 3 and the UNI at SD-WAN Edge 4 require the concatenation of multiple TVCs, A–F, C–G, or D–G. This requires routing through another SD-WAN Edge. The SD-WAN Service Provider is responsible for connectivity between SD-WAN Edges that are attached to different Underlay Connectivity Services. How this is done is beyond the scope of this document.

TVCs do not have Service Attributes that can be used by the Subscriber in Policy definitions, but they have several important characteristics:

- A TVC appears to each SD-WAN Edge as a single routing hop.
- If a TVC traverses an Internet Access UCS, then the TVC is Public; otherwise, it is Private.
- The SD-WAN Edge can implement encryption on a TVC basis, in which case all traffic forwarded across the TVC is encrypted.
- The performance of a TVC is one of the factors that affects whether the TVC is used for forwarding each Application Flow.

TVCs are the responsibility of the Service Provider and are not agreed on with the Subscriber. How they are instantiated, when they are instantiated, and whether they are ephemeral or long-lasting are all beyond the scope of this document. As noted above, there are no Service Attributes for TVCs, and this specification does not provide a mechanism for the Subscriber to refer to specific TVCs in Policy definitions. The Service Provider may or may not choose to expose information about TVCs to the Subscriber.

7.11 Application Flow, Application Flow Specification, and Policy

A defining characteristic of an SD-WAN Service is the forwarding of IP Packets over an SWVC, where the SWVC may operate over multiple UCSs. UCSs may have different Service Attributes and characteristics. Policies are used to forward Application Flows over the most appropriate UCS(s).

An Application Flow is a sequence of IP Packets that ingress at an SD-WAN UNI or are directed towards an SD-WAN UNI from a UCS that:

1. Have the same direction (ingress or egress), and
2. Match the same Application Flow Specification, and
3. Are in the same Zone (i.e., have source IP Addresses in the same Zone or are all in the Zone *Internet*)

An Ingress Application Flow is a subset of the Ingress IP Packets at a UNI, and an Egress Application Flow is a subset of the IP Packets destined to egress at a UNI. The description is slightly different between ingress and egress because IP Packets in an Ingress Application Flow always arrive at the Ingress UNI, but IP Packets in an Egress Application Flow might be discarded by a Policy and therefore won't reach the Egress UNI.

Application Flows contain distinct sets of IP Packets; in other words, each IP Packet belongs to at most one Ingress Application Flow and at most one Egress Application Flow.

Policies are assigned to each Application Flow at an SWVC End Point. An Ingress Policy describes the required forwarding behavior for Ingress Application Flows (e.g., should they be forwarded or discarded, should they be encrypted, should they be rate limited, and several more). An Egress Policy describes whether Egress Application Flows directed toward a UNI should be forwarded to that UNI or discarded due to security or other considerations.

Section 8 provides more detail about Application Flows and Policies.

7.12 Zone

An IP host is either a physical or logical IP addressable network interface. The IP hosts in the Subscriber Network can be separated into groups that have different business and security requirements. Zones are used to define and describe these groups of IP hosts. At least one Zone is required in every SWVC. An SD-WAN Service can be used to forward traffic for some or all these Zones. Some examples of Zones are:

- Corporate Zone vs. Guest Zone
- Point-of-sale Zone vs. Corporate Zone
- Engineering Zone vs. Finance Zone vs. HR Zone
- Company Zone vs. Guest Zone vs. Building Management System Zone (or IoT Zone)

Note that the Zone *Internet* is reserved. See section 9.5.

An Application Flow is defined by the UNI, the Application Flow Specification that it matches, and its Zone (i.e., set of source IP hosts). Since Policies are assigned to Application Flows, IP Packets at a given UNI that match the same Application Flow Specification can have different treatment if they are in different Zones.

Terminology Note: This document uses the phrase “IP Packet in Zone *xxx*” to indicate that at the Ingress UNI, the IP Packet was received from an IP host that is in Zone *xxx*. In addition, the phrase “IP Packet in Zone *Internet*” indicates that the IP Packet entered the SWVC directly from the Internet, not via a UNI.

The SWVC List of Zones Service Attribute (section 9.5) specifies the Zones used by the SD-WAN Service.

One of the goals of using Zones is commonality in handling Application Flows at different locations in the SD-WAN Service. For example, the same Policies can be assigned to Application Flows in “Guest-Wifi” Zone at every SWVC End Point in the SD-WAN Service that has a “Guest-Wifi” Zone. Hence, the SWVC List of Zones Service Attribute includes, for each Zone, a set of

Zone-wide Application Flow-to-Policy mappings. Unless overridden by another Policy assigned at an SWVC End Point, these Policies are assigned at every SWVC End Point where the associated UNI has the Zone defined (see section 8.4).

A Zone definition includes an SWVC-wide list of IP prefixes that defines the IP subnets that compose the Zone. A host is assigned to a Zone if the IP address of the host is a member of one of the IP Prefixes for the Zone. If the source of an IP Packet or the destination of an IP Packet is not assigned to a Zone, the IP Packet is discarded (see [R23]).

Zones can reduce the number of Application Flow Specifications but may increase the number of distinct Application Flows identified at UNIs in the SWVC. A single Application Flow Specification, for example one called *email*, can be matched at an SD-WAN UNI against IP Packets in (arriving from) two Zones, e.g., *HR* and *Finance*, resulting in two Application Flows at the UNI, $\langle email, HR \rangle$ and $\langle email, Finance \rangle$, and each one can have a different Policy assigned.

7.13 Virtual Topology

For many Application Flows, it might be desirable for IP Packets arriving at any UNI to be able to be forwarded to any other UNI, but for other Application Flows, the desired forwarding behavior might be more limited. As an example, consider a retail subscriber with multiple stores. For a person-to-person messaging Application Flow across the company, a full mesh is useful, but the point-of-sale Application Flow in each store can only forward IP Packets to the POS controller, at the company data center.




Figure 7 – Different Virtual Topologies in an SWVC

The Application Flows in Figure 7 experience two different forwarding behaviors in the same SD-WAN Service.

There can be multiple forwarding topologies in the SWVC. This is achieved by defining Virtual Topologies (see section 9.6), each of which includes a subset of the SWVC End Points (the full mesh of all the SWVC End Points is one of the possible Virtual Topologies). Each Application Flow is assigned to a Virtual Topology by Policy (see section 9.12.2.4).

7.14 Internet Breakout

In general, an SD-WAN Service delivers IP Packets in an Ingress Application Flow from the Ingress SD-WAN UNI to another (Egress) SD-WAN UNI, over one or more TVCs. If some of the Underlay Connectivity Services for the SD-WAN Service are Internet Access Services, then some those TVCs may traverse the Internet.

An SD-WAN Edge that is connected to an Internet Access UCS can also forward IP Packets in an Application Flow directly to the Internet, rather than over a TVC to another SD-WAN UNI. Thus, the SD-WAN Service can act as an “Internet gateway” for traffic from the Subscriber Network to IP hosts on the Internet.

Whether IP Packets in an Application Flow are delivered to the Internet by the SD-WAN Service (i.e., via one of the SD-WAN Edges) is controlled by the INTERNET-BREAKOUT Policy Criterion in the Policy assigned to the Application Flow, and by routing information within the SD-WAN Service. If INTERNET-BREAKOUT is *Disabled*, then IP Packets in the Application Flow can only be delivered to SD-WAN UNIs (i.e., they cannot “break out” to the Internet). If INTERNET-BREAKOUT is *Enabled*, then each IP Packet in the Application Flow is delivered either to an SD-WAN UNI or to the Internet via an Internet Access UCS, depending on whether the destination IP Address in the IP Packet is reachable via an SD-WAN UNI or via the Internet.

An example of Internet Breakout is shown in Figure 8 (this is the same diagram as Figure 3 with the addition of an Internet-connected site outside of the SD-WAN Service). It shows an Ingress IP Packet at site B being forwarded across the UCS UNI for UCS #1 (an Internet Access UCS), but instead of being sent over one of the TVCs, it is forwarded to an Internet destination.

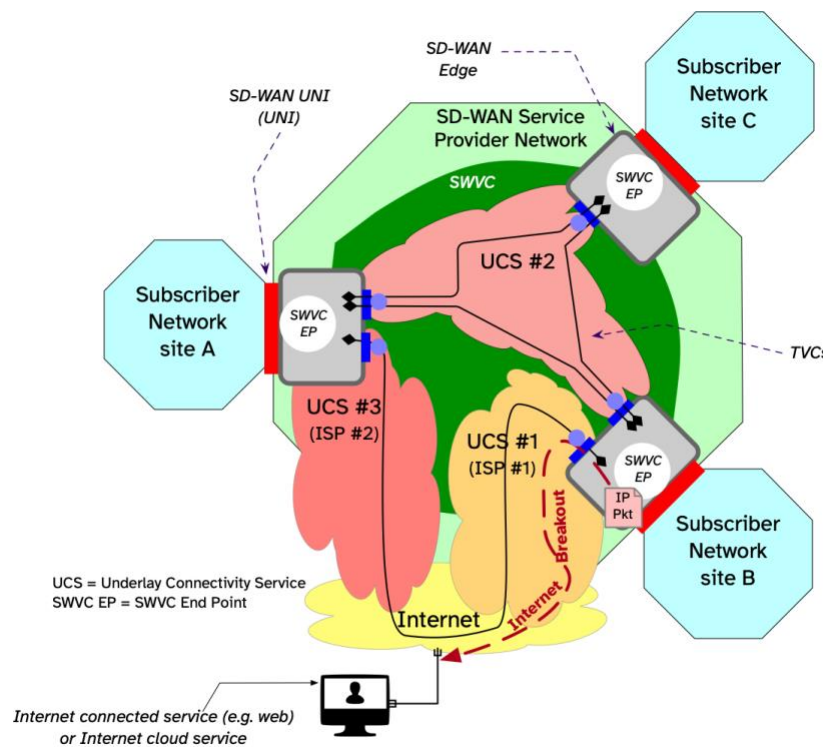


Figure 8 – Internet Breakout

If IP Packets destined to Internet hosts are forwarded via an Internet Access UCS that is connected to the SD-WAN Edge where the Ingress UNI for those IP Packets is located (as shown in Figure 8), this is called Local Internet Breakout. If an appropriate Internet Access UCS is not available at the SD-WAN Edge containing the Ingress UNI for the Application Flow, the Service Provider delivers the IP Packet over the SD-WAN Service (i.e., over one or more TVCs) to another SD-WAN Edge for “breakout” to the Internet. This would be the case if:

- There is no Internet Access UCS at the (ingress) SD-WAN Edge (e.g., for IP Packets that ingress at site C),
- There is an Internet Access UCS at the (ingress) SD-WAN Edge, but it is currently inaccessible/down,
- There is an Internet Access UCS at the SD-WAN Edge, but it is not enabled for breakout. This is controlled by the UCS End Point Breakout Service Attribute (see section 12.3.3).

Note that disabling Internet Breakout for an Application Flow does not preclude the SD-WAN Service from delivering IP Packets in the Application Flow to another UNI via a TVC that traverses an Internet Access UCS.

If at least one Application Flow at a UNI has a Policy that indicates Internet Breakout is *Enabled*, then IP Packets received from Internet sources (rather than from another UNI) are allowed to egress at that UNI (subject to any Egress Policies applied). Otherwise, IP Packets received from Internet sources cannot be delivered to that UNI.

7.15 Rate Limiter and Bandwidth Flow

It is frequently desirable to limit the rate and/or commit to a minimum rate that data for an Application Flow (or a set of Application Flows) can enter the SD-WAN network. This is achieved by including a BANDWIDTH Policy Criterion (see section 9.12.2.10) in the Policy assigned to the Application Flow.

For example, it may be desirable to:

1. commit to a minimum of 100 Mb/s and a limit of 200 Mb/s for the “Chat” Application Flow Specification in Zone “Corporate” at a UNI, and
2. commit to a minimum of 75 Mb/s and a limit of 150 Mb/s for the “Internet” Application Flow Specification in Zone “Corporate” at the UNI.

In addition, it may be desirable to limit the rate and/or commit to a minimum rate for a set of Application Flows in aggregate. For example, it may be desirable to:

3. commit to a minimum of 250 Mb/s and a limit of 500 Mb/s for the set of all Application Flows in the Zone “Guest Access” at the UNI.

The difference is that in bullets 1 and 2, in Zone “Corporate”, each Application Flow has its own individual bandwidth constraints and thus each Application Flow is unaffected by other Application Flows. In contrast, in bullet 3, in Zone “Guest Access” all the Application Flows share the same bandwidth constraints, and thus each Application Flow can be affected by the behavior of other Application Flows in the set.

The term Bandwidth Flow is defined as a set of one or more Application Flows that are treated as a single IP packet flow for the purpose of enforcing bandwidth constraints (commitments and limits). In the previous example, at the UNI, there are three Bandwidth Flows that have been mentioned:

- Application Flow Specification “Chat” in Zone “Corporate”
- Application Flow Specification “Internet” in Zone “Corporate”
- Application Flow Specifications {“Chat”, “Internet”, etc.} in Zone “Guest Access”

Commitments and limits for Bandwidth Flows are specified using Rate Limiters. A Rate Limiter is an abstract component of an SD-WAN Service that ensures that the information rate (measured over time periods of duration *irduration*, see section 9.7) of a Bandwidth Flow conforms to the parameters of the Rate Limiter. A Rate Limiter meters a Bandwidth Flow and, if necessary, delays or discards IP Packets in the Bandwidth Flow so that it conforms to the Rate Limiter parameters. The effect of metering a Bandwidth Flow—that is, comparing the actual sequence of IP Packets that pass the metering point to the description in terms of the Rate Limiter parameters—is to declare IP Packets in the Bandwidth Flow either conformant or nonconformant. IP Packets that are not conformant are discarded in order to ensure that the Bandwidth Flow meets the specified bandwidth limits and commitments. A Rate Limiter commonly behaves as a traffic policer and may also include traffic shaping.

Rate Limiters are specified using a two-tuple, $\langle commit, limit \rangle$ where:

- *commit* – is the threshold information rate (bits per second) at or below which the SD-WAN Service Provider commits to deliver packets in the Bandwidth Flow with high probability under all traffic conditions, i.e., regardless of the information rate for other Bandwidth Flows at this UNI, or at other UNIs in the same SD-WAN Edge. The value for *commit* can be 0.
- *limit* – is the threshold information rate (bits per second) above which the Service Provider does not deliver IP Packets in the Bandwidth Flow under any traffic conditions, i.e., the SD-WAN Service Provider delays or discards IP Packets in the Bandwidth Flow regardless of the bandwidth used for other Bandwidth Flows at this UNI, or at other UNIs in the same SD-WAN Edge. The value for *limit* can be *None*.

[R7] In the specification of a Rate Limiter, the value of the *limit* element, if not *None*, **MUST** be greater than or equal to the value of the *commit* element.

A Rate Limiter can be unnamed or named.

An unnamed Rate Limiter is specified by including a BANDWIDTH Policy Criterion specifying the *commit* and *limit* values for the Rate Limiter (section 9.12.2.10) in the Policy assigned to an Application Flow. An unnamed Rate Limiter applies only to that Application Flow. Therefore, an Application Flow that is assigned a Policy that includes a BANDWIDTH Policy Criterion with an unnamed Rate Limiter is always the only Application Flow in its Bandwidth Flow.

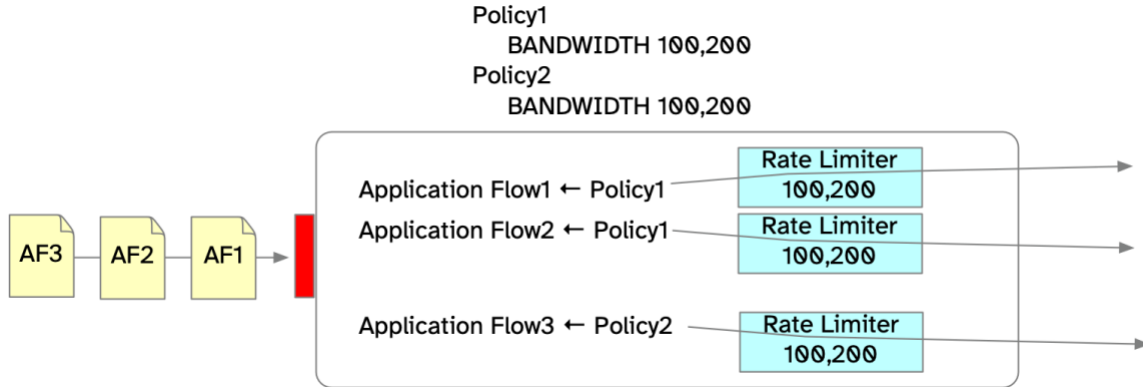


Figure 9 – Example of Unnamed Rate Limiters

In Figure 9, Policy1 and Policy2 both include a BANDWIDTH Policy Criterion with an unnamed Rate Limiter. Policy1 is assigned to Application Flow 1 and Application Flow 2 at the SWVC End Point, and Policy2 is assigned to Application Flow 3. Each of the three Application Flows is independently rate limited. In other words, each of these Application Flows is a separate Bandwidth Flow even though the Rate Limiters parameters are the same in both Policies.

A named Rate Limiter is defined in the SWVC List of Rate Limiters Service Attribute (see section 9.9). Each entry in that list specifies a Rate Limiter 3-tuple <name, commit, limit>.

If one or more Policies include a BANDWIDTH Policy Criterion that specifies the name of the same named Rate Limiter (rather than explicit *commit* and *limit* values), only one instance of that Rate Limiter is created at the SWVC End Point and all Application Flows that are assigned any Policy that includes a BANDWIDTH Policy Criterion specifying that named Rate Limiter share the Rate Limiter. Therefore, all the Application Flows that are assigned any Policy that includes a BANDWIDTH Policy Criterion specifying the same named Rate Limiter are members of the same Bandwidth Flow.

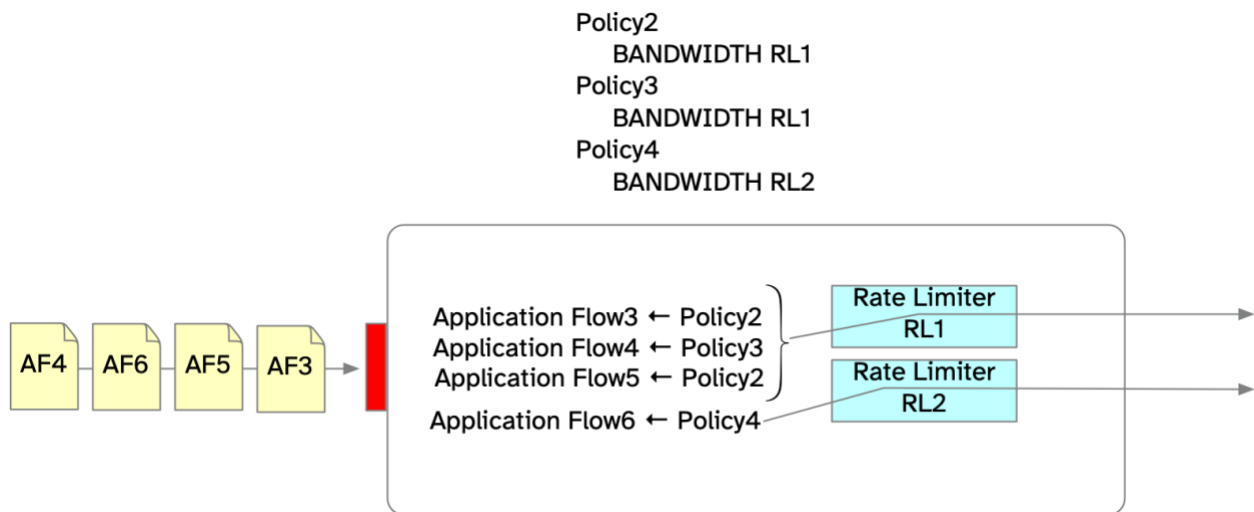


Figure 10 – Example of Named Rate Limiters

In Figure 10 three Policies (Policy 2, 3 and 4) are defined. Policies 2 and 3 include a BANDWIDTH Policy Criterion with a named Rate Limiter, RL1, and Policy 4 is defined with a different named Rate Limiter, RL2. Each Policy is assigned to Application Flow(s) at the SWVC End Point. Application Flows 3, 4, and 5 share RL1 since they are all assigned a Policy that specifies RL1. RL1 views those three Application Flows as a single Bandwidth Flow. Application Flow 6 currently has RL2 to itself, but subsequent Policy assignments (not shown in Figure 10) referencing that named Rate Limiter would result in other Application Flows sharing that Rate Limiter.

Specifying a value greater than 0 for *commit* indicates that delivering at least this rate is necessary to achieve the desired or intended Bandwidth Flow behavior, whereas *commit* = 0 indicates that the Bandwidth Flow has no bandwidth requirement and operates with whatever bandwidth is available.

The intent of specifying a value other than *None* for *limit* is to indicate an upper bound on the bandwidth used by the Bandwidth Flow. The *limit* can be used, for example, to ensure that there is sufficient bandwidth for other Bandwidth Flows. Specifying *limit* = *None* means that there is no maximum imposed on the Bandwidth Flow up to the limits imposed by the SD-WAN UNI speed (which needs to be agreed on between the Subscriber and Service Provider) and Underlay Connectivity Service bandwidth constraints.

The intended behavior is shown in the following diagram.

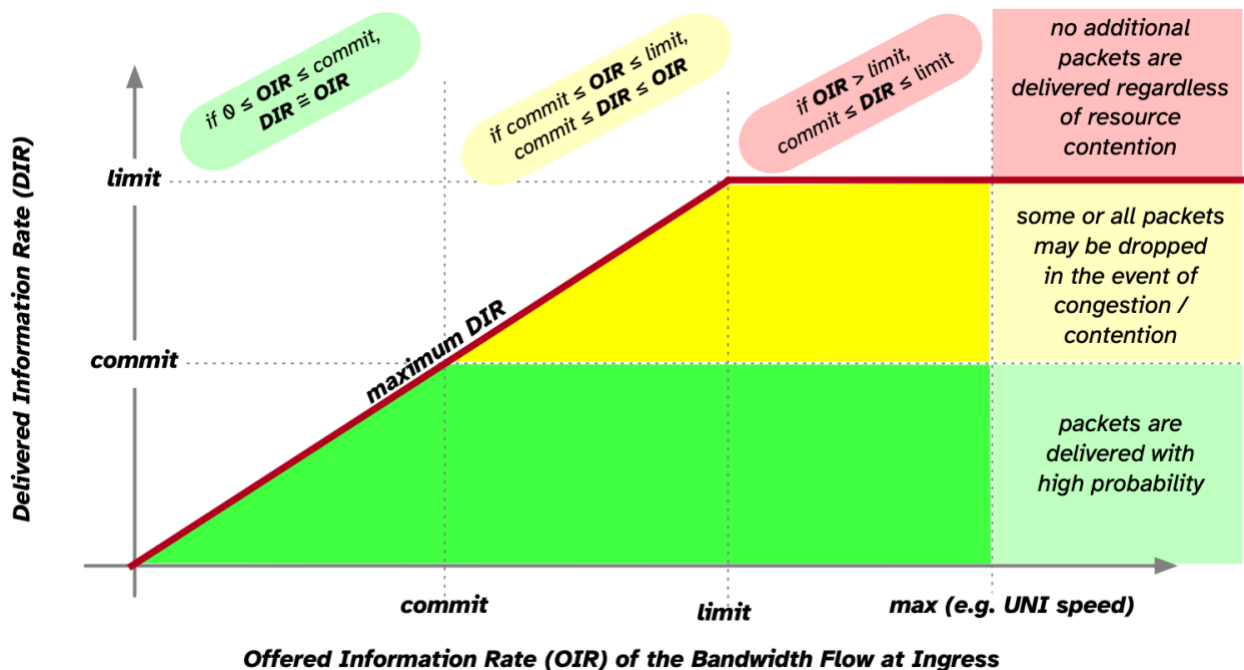


Figure 11 – Intended Behavior of a Rate Limiter

For an offered information rate from zero up to the *commit* bandwidth, the expectation is that IP Packets in the Bandwidth Flow are delivered with high probability. In the event of network congestion in the forwarding Path, such packets in the Bandwidth Flow have a low probability of

discard or delay. They could be discarded or delayed due to other Policy Criteria or by Security Policies, but not due to bandwidth limiting.

For an offered information rate between *commit* and *limit*, IP Packets in the Bandwidth Flow are delivered with high probability unless there is contention with IP Packets in other Bandwidth Flows for the available UCS bandwidth. In this case, the Service Provider might delay or discard IP Packets in this Bandwidth Flow to reduce its impact on UCS utilization, but not such that the delivered information rate is reduced below *commit*.

For an offered information rate above *limit*, no additional packets will be delivered beyond those delivered for an offered information rate equal to *limit*.

Note that if the value of *commit* is 0, the yellow area would extend down to the horizontal axis (covering the green area) and if *limit=None*, the yellow area would extend upwards to the limit imposed by the UNI.

7.16 SD-WAN IP Packet Delivery

An SD-WAN Service delivers IP Packets between Subscriber Locations. In that sense, it shares many attributes with a MEF IP Service (as described in MEF 61.1, *IP Service Attributes* [32]). Therefore, many of the sections of this document are derived from MEF 61.1.⁵

Since IP Packet forwarding in SD-WAN Services is based on Application Flows and Policies, there are several Service Attributes defined to describe these capabilities, and several of the Service Attributes integrated from MEF 61.1 have been modified to focus on Application Flow-based forwarding rather than general IP/layer 3-based forwarding. An IP Service can be used as an Underlay Connectivity Service for an SD-WAN Service.

An SD-WAN Service (SWVC) provides the logical construct of an IP Virtual Private Routed Network for a Subscriber. This section describes the basic IP forwarding paradigm for an SD-WAN Service, and requirements that indicate which fields of an IP Packet can be modified while traversing the SD-WAN Service and under what conditions they can be modified.

7.16.1 IP Packet Forwarding

The basic forwarding paradigm for an SD-WAN Service is a Virtual Private Routed Network (VPRN) as described in RFC 2764 [13]. As a VPRN, the SD-WAN Service relies on destination-based IP forwarding (based on standard IP longest prefix match), which can then have additional forwarding constraints as a result of Policies applied by the SD-WAN Service.

The Subscriber and the Service Provider may use static routing or a dynamic routing protocol at the UNI to exchange information about reachable IP Prefixes.

[R8] The Service Provider **MUST NOT** deliver an ingress IP Packet to a UNI where the destination address is not reachable.

⁵ This is to ensure the greatest level of commonality between the two specifications as well as for expediency.

SD-WAN Services forward unicast, multicast, and broadcast IP Packets; however, this specification includes requirements for forwarding unicast IP Packets only. Requirements for forwarding of multicast and broadcast IP Packets are out of scope for this version of the standard.

The routing context of the SD-WAN Service may be independent of any routing contexts for the Underlay Connectivity Services used by the SD-WAN Service.

7.16.2 IP Packet Transparency

In general, an SWVC conveys IP Packets without modifying the contents; however, there are some exceptions that are captured in the following requirements:

- [R9]** If an Ingress IPv4 Packet is mapped to an SWVC and delivered as an Egress IPv4 Packet, and the packet has not been fragmented as described in RFC 791 [6], the Egress IPv4 Packet **MUST** be identical to the Ingress IPv4 Packet except that the following fields in the IPv4 header can be changed:
- The TTL field (RFC 791 [6])
 - The DS (RFC 3260 [17]) and ECN (RFC 3168 [16]) fields
 - The Loose Source and Record Route option, the Strict Source and Record Route option, and the Record Route option, if present in the packet (RFC 791 [6])
 - The Destination Address field, if the Loose Source and Record Route option or the Strict Source and Record Route option are present in the packet (RFC 791 [6])
 - The Header Checksum field (RFC 791 [6])
 - Any other field(s), subject to agreement between the Subscriber and the Service Provider
- [R10]** If an Ingress IPv4 Packet is mapped to an SWVC and is fragmented by the Service Provider as described in RFC 791 [6] resulting in a number of corresponding IPv4 Packets that are delivered as Egress IPv4 Packets, the Egress IPv4 Packets **MUST** be such that reassembly as described in RFC 791 [6] results in an IP Packet that is identical to the Ingress IPv4 Packet except that the following fields in the IPv4 header can be changed:
- The TTL field (RFC 791 [6])
 - The DS (RFC 3260 [17]) and ECN (RFC 3168 [16]) fields
 - The Loose Source and Record Route option, the Strict Source and Record Route option, and the Record Route option, if present in the packet (RFC 791 [6])
 - The Destination Address field, if the Loose Source and Record Route option or the Strict Source and Record Route option are present in the packet (RFC 791 [6])
 - The Header Checksum field (RFC 791 [6])
 - Any other field(s), subject to agreement between the Subscriber and the Service Provider

- [R11]** If an Ingress IPv6 Packet is mapped to an SWVC and delivered as an Egress IPv6 Packet, the Egress IPv6 Packet **MUST** be identical to the Ingress IPv6 Packet except that the following fields in the IPv6 header can be changed:
- The Hop Limit field (RFC 8200 [28])
 - The DS (RFC 3260 [17]) and ECN (RFC 3168 [16]) fields
 - The value of any options within a Hop-by-Hop Options header (if present) that have the third high-order bit in the option type field set (RFC 8200 [28])
 - Any other field(s), subject to agreement between the Subscriber and the Service Provider

The use of the Loose Source and Record Route option, the Strict Source and Record Route option, and the Record Route option in IPv4 packets can cause problems due to the additional processing needed at each hop along the path. In addition, the Loose Source and Record Route option and the Strict Source and Record Route option open a number of potential security risks as documented in RFC 6274 [25], which outweigh any legitimate use.

- [D1]** A Service Provider **SHOULD** discard Ingress IPv4 Packets that contain the Loose Source and Record Route option, the Strict Source and Record Route option, or the Record Route option.

7.17 Identifier String

Many of the Service Attribute values in this document are strings that are used for identification of an element. The document uses a single definition for the structure of these Identifier Strings.

The length of the Identifier String is not limited; however, since it is used in human interfaces, it is best to avoid very long Identifier Strings.⁶

Since the Identifier String is used in human interfaces, the allowable character set is chosen to contain printable characters.

- [R12]** An Identifier String **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.

Identifier Strings only include printable ASCII characters.

⁶ MEF 61.1 [32] limits these strings to 53 octets for MEF IP Services.

8 Application Flows and Policies

In an SD-WAN Service, IP Packets are classified into “Application Flows” at the Ingress and Egress UNIs and a “Policy” is assigned to each Application Flow to describe how the Service handles the Application Flow. There are several important logical constructs used in this standard to describe and define Application Flows and Policies:

- **Application Flow**
 - Application Flow Specification
 - Application Flow Criterion
 - Application Flow Specification Group
- **Policy**
 - Policy Criterion

There are several Service Attributes that relate to assigning Policies to Application Flows:

- SWVC List of Application Flow Specifications Service Attribute (section 9.14)
- SWVC List of Application Flow Specification Groups Service Attribute (9.13)
- SWVC List of Policies Service Attribute (9.12)
- SWVC List of Zones Service Attribute (section 9.5)
- SWVC End Point Policy Map Service Attribute (10.4)

The following subsections describe these constructs in detail and their relationship to each other.

8.1 Application Flows

An Application Flow is the subset of the IP Packets that ingress at a UNI or are destined for egress at a UNI that match a specific Application Flow Specification and are in a specific Zone.

The description is slightly different between ingress and egress because IP Packets in an Ingress Application Flow always arrive at the Ingress UNI, but IP Packets in an Egress Application Flow might be discarded by a Policy and therefore will not reach the Egress UNI. An Ingress IP Packet can be mapped to, at most, one Ingress Application Flow at a UNI and can belong to, at most, one Egress Application Flow. The Ingress IP Packets that are mapped to a given Ingress Application Flow at a UNI will not necessarily belong to the same Egress Application Flow since different Ingress IP Packets can be destined to different UNIs.

Although the term “Application Flow” includes the word “Application” there is only a loose connection between the two. An Application Flow can map to (what is typically thought of as) an application, but it can also represent a subset or a superset. For example:

- An Application Flow can include IP Packets for several individual computer applications, such as “all packets that use the RTP protocol as defined in RFC 3550 [19]”, or
- IP Packets for a single application could be split among multiple Application Flows, such as a single video conferencing call resulting in a “Video” Application Flow and an “Audio” Application Flow, or an IP phone resulting in an “Audio” Application Flow and a “Control/Signaling” Application Flow, or

- An Application Flow can include all IP Packets from an IP address range, such as 10.10.10.x/24, which could, for example, represent all Point-of-Sale terminals at a location.

In this standard, Policies are assigned to Application Flows, not applications. As part of the SD-WAN Service, the Subscriber and the Service Provider agree on the specifications for Application Flows that are identified at the SD-WAN UNIs. In the IP phone example in the second bullet above, the SD-WAN Service's handling of the audio channel can be very different than its handling of the signaling channel, because a different Policy can be assigned to each. It is also possible to define a single Application Flow in such a way that it would include both of those components, in which case there would only be a single Policy assigned.

8.2 Application Flow Specifications and Application Flow Criteria

An "Application Flow Specification" describes a set of characteristics of IP Packets that are used to assign the IP Packets to Application Flows. These characteristics, referred to as "Application Flow Criteria", can include matching specific packet fields (in any layer) against specific values, and can also include more complex algorithms and heuristics that inspect sequences of IP Packets. The SWVC List of Application Flow Specifications Service Attribute (section 9.14) describes the Application Flow Specifications that are available for the SWVC and the criteria (Application Flow Criteria) that compose them. Application Flow Specifications are used to categorize IP Packets that ingress at a UNI as well as IP Packets that are destined to an Egress UNI. The Service Provider is required to support some of the Application Flow Criteria specified in this standard (see Table 8), and support is recommended for the other Application Flow Criteria (see Table 9). In many cases, the Service Provider provides a catalog of named, pre-defined matching criteria (accessible via the APPID Application Flow Criterion specified in section 9.14), and the Subscriber can select from the catalog. The Service Provider's catalog includes a description of the IP Packets that are matched by each of the items in the catalog.

It is often desirable to assign different Policies to IP Packets based on the Zone of the source IP host. For example, an imap (email) IP Packet from Zone *engineering* may require different handling than the same IP Packet from Zone *corporate*. This concept is shown (for ingress) in Figure 12. Clearly, if the Subscriber has a single Zone, then, in effect, only the Application Flow Specification is relevant.

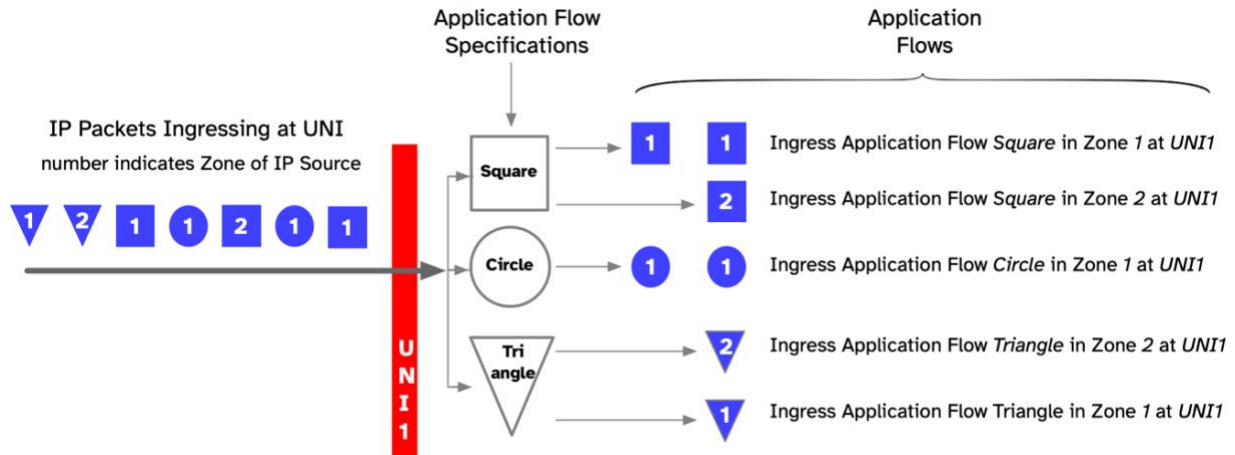


Figure 12 – Examples of Application Flows⁷

Figure 12 depicts a stream of IP Packets arriving at UNI1 and being classified into Application Flows based on their Application Flow Specification and Zone. For example, two of the Ingress IP Packets are classified into Ingress Application Flow *Square* in Zone 1 at UNI1 whereas only one is classified into Ingress Application Flow *Square* in Zone 2 at UNI1. Similarly, two of the Ingress IP Packets are classified into Ingress Application Flow *Circle* in Zone 1 at UNI1.

A reference⁸ to an Application Flow includes:

- the UNI,
- the direction (ingress or egress),
- the name of the Application Flow Specification that it matches, and
- the name of the Zone that it matches

Some of these may be understood contextually. The description of each of the 5 Application Flows on the right side of Figure 12 includes all four of these components.

8.3 Application Flow Specification Groups

Each Application Flow Specification can be a member of an “Application Flow Specification Group”. An Application Flow Specification Group provides a way to assign the same Policy to multiple Application Flow Specifications with a single assignment. See section 8.4 for more details. The SWVC List of Application Flow Specification Groups Service Attribute is described in section 9.13.

8.4 Policies and Policy Criteria

A Policy is a named list of Policy Criteria that can be assigned to an Application Flow and that determines how an SD-WAN Service handles IP Packets in the Application Flow. Each Policy Criterion describes a particular aspect of the handling of IP Packets to which the Policy is assigned.

⁷ This diagram is not intended to imply any particular ordering of operations. An implementation could match Zone before Application Flow Specification, or after, or in parallel.

⁸ There is no need in this document to formally name Application Flows, for example by assigning identifiers, but there is often a need to refer to them in explanatory text.

These aspects include requirements on UCS and TVC characteristics, forwarding requirements, security requirements, bandwidth commitments and limitations, performance goals, etc.

Policies are assigned to Application Flows at each SWVC End Point. Ingress Policies specify the desired behavior concerning forwarding of Ingress Application Flows and Egress Policies specify whether IP Packets in Egress Application Flows should be transmitted at the UNI or discarded based on security or other considerations.

There are two Service Attributes that can assign a Policy to an Application Flow:

- SWVC List of Zones Service Attribute
- SWVC End Point Policy Map Service Attribute

A Zone definition in the SWVC List of Zones Service Attribute (see section 9.5) can include a Zone-wide Ingress Policy for Ingress IP Packets that match a specified Application Flow Specification or the Application Flow Specifications in an Application Flow Specification Group. This, in effect, becomes the default Ingress Policy at all SWVC End Points for Application Flows that match the Application Flow Specification(s) and that Zone, unless overridden.

The SWVC End Point Policy Map Service Attribute (section 10.4) allows an Ingress Policy and an Egress Policy to be assigned to a $\langle \text{Zone}, \text{Application Flow Specification} \rangle$ at the UNI for the SWVC End Point (i.e., an Application Flow) or to a $\langle \text{Zone}, \text{Application Flow Specification Group} \rangle$ at that UNI (one or more Application Flows).

This results in four ways that a Policy can be assigned to an Ingress Application Flow:

- Zone-wide Policy for an Application Flow Specification
- Zone-wide Policy for an Application Flow Specification Group
- SWVC End Point Policy Map for Application Flow Specification
- SWVC End Point Policy Map for Application Flow Specification Group

The following diagram shows many of the possible cases.

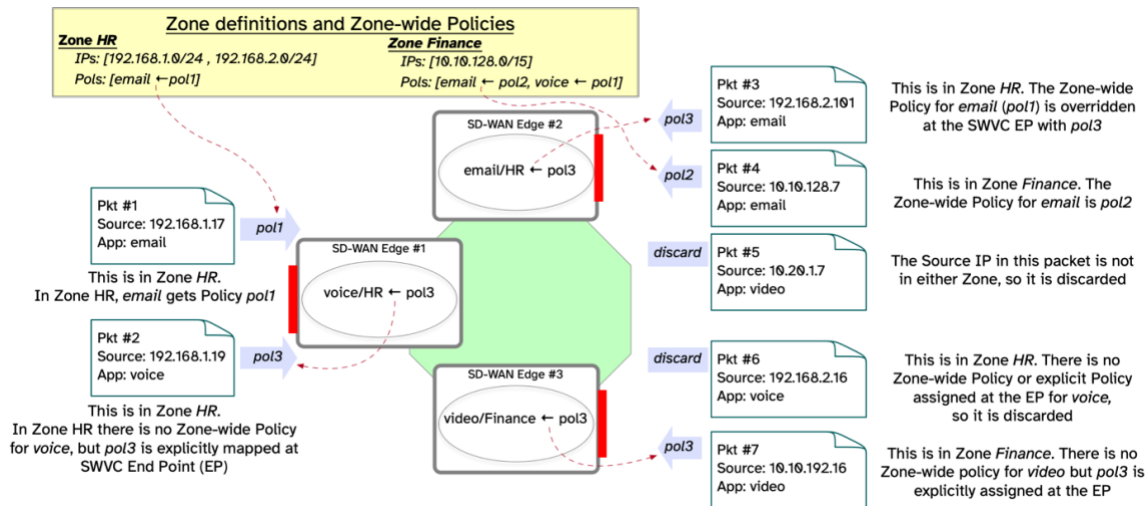


Figure 13 – Zones and Policy Assignment Examples

Figure 13 depicts several examples of the options for Policy assignment. There are two Zones, *HR* and *Finance*. *HR* has two IP subnets and one Zone-wide Policy (for *email*). *Finance* has one IP subnet and two Zone-wide Policies (one for *email* and one for *voice*). Note that Policy *pol1* is specified as a Zone-wide policy for different Application Flows in the two Zones. At each SWVC End Point Policy (inside the oval), *pol3* is explicitly mapped to one of the Application Flow/Zone combinations. Seven packets are presented at the three UNIs. The fat arrow outside of each packet indicates the Policy that is assigned to the packet at that SWVC End Point. Below or next to each packet is an explanation of why the Policy was assigned.

Since there are multiple mechanisms for assigning a Policy to a single Ingress Application Flow, there is a well-defined precedence as described in section 10.4.1.

For Egress Application Flows, the SWVC End Point Policy Map Service Attribute (see section 10.4) can associate an Egress Policy to an Application Flow. As with Ingress Application Flows, Egress Application Flows at an SWVC End Point (associated with the Egress UNI) are identified by matching an Application Flow Specification and a Zone, with the special case that IP Packets arriving directly from the Internet (i.e., not over a TVC) are assigned the reserved Zone name *Internet*. Section 10.4.2 provides additional detail regarding the assignment of Policies to Egress Application Flows.

When a Policy is associated with an Application Flow Specification Group in a Zone at an SWVC End Point, the Policy is applied to each Application Flow that matches an Application Flow Specification that is a member of the Application Flow Specification Group. For example, using the Application Flow Specifications in Figure 12, *Square*, *Circle*, and *Triangle* could all be members of an Application Flow Specification Group *shapes*. A Policy associated with the Application Flow Specification Group *shapes* in Zone 1 at UNI1 is assigned to the three resulting Application Flows *uni1/zone1/Square*, *uni1/zone1/Circle*, and *uni1/zone1/Triangle*. Each of the Application Flows can, however, have an explicit Policy assignment that overrides the Policy assigned to the Application Flow Specification Group (see section 10.4.1).

9 SD-WAN Virtual Connection (SWVC) Service Attributes

This section contains Service Attributes that apply to an SD-WAN Virtual Connection. There is one instance of these attributes for each SD-WAN Virtual Connection. The attributes are summarized in the following table, and each is described in more detail in the subsequent sections.

Attribute Name	Summary Description	Possible Values
SWVC Identifier	Identification of the SWVC for management purposes	Unique Identifier String for the SD-WAN Service.
SWVC List of SWVC End Points	The SWVC End Points that are associated by the SWVC	List of SWVC End Point Identifiers
SWVC List of UCSs	The Underlay Connectivity Services that are used by the SWVC	List of UCS Identifiers
SWVC Reserved Prefixes	IP Prefixes reserved for use by the SP	List of IP Prefixes
SWVC List of Zones	A list of the Zones supported at one or more UNIs in the SD-WAN Service and Zone-wide Policies associated with the Zones	List of 3-tuples <name, list of IP Prefixes, list of policy mappings>
SWVC List of Virtual Topologies	A list of named Virtual Topologies that can be assigned to Application Flows by Policy	List of 3-tuples (vtName, vtType, vtEP)
SWVC Performance Time Intervals	Time intervals used in the evaluation of the PERFORMANCE and BANDWIDTH Policy Criteria and the Performance Metrics	3-tuple (evalinterval, arrivalinterval, irduration)
SWVC Service Objectives Reporting Periods	Definition of reporting periods for evaluating Service Performance Objectives.	2-tuple (s, T)
SWVC List of Rate Limiters	A list of named Rate Limiters that can be assigned to Application Flows by Policy	List of 3-tuples (name, commit, limit)
SWVC List of Security Policies	A list of the Security Policies (defined in MEF 88) that are available for use in this SWVC	List of Security Policies (Identifier Strings)
SWVC List of Application Flow Business Importance Levels	An ordered list of labels that can be used in Policies to indicate the relative business importance of an Application Flow	None or 2-tuple (list of labels, default label)
SWVC List of Policies	A list of the Policies that can be applied to Application Flows carried by the SWVC	List of 2-tuples (Policy Name, List of Policy Criteria n-tuples)
SWVC List of Application Flow Specification Groups	A list of the Application Flow Specification Groups. Each Application Flow Specification can, optionally, be a member of one of the Application Flow Specification Groups.	List of Application Flow Specification Group Names (Identifier Strings)
SWVC List of Application Flow Specifications	A list of the Application Flow Specifications that are recognized by the SD-WAN Service	List of 3-tuples (Application Flow Specification Name, List of Application Flow Specification Criteria, Application Flow Specification Group Name)

Table 4 – Summary of SWVC Service Attributes

9.1 SWVC Identifier Service Attribute

The value of the SWVC Identifier Service Attribute is a string that is used by the Subscriber and the Service Provider to uniquely identify an SWVC.

- [R13] The value of the SWVC Identifier Service Attribute **MUST** be an Identifier String.
- [R14] The value of the SWVC Identifier Service Attribute **MUST** be unique across all SWVCs in the Service Provider Network.

9.2 SWVC List of SWVC End Points Service Attribute

The value of the SWVC List of SWVC End Points Service Attribute is a non-empty list of SWVC End Point Identifier Service Attribute values (section 10.1). The list contains one SWVC End Point Identifier value for each SWVC End Point connected by the SWVC.

- [R15] An SWVC End Point Identifier **MUST NOT** appear more than once in the value of the SWVC List of SWVC End Points Service Attribute.
- [R16] The SWVC End Points in the value of the SWVC List of SWVC End Points Service Attribute **MUST** all be associated with different UNIs.
- [R17] If an Egress IP Packet at an SWVC End Point results from an Ingress IP Packet at a different SWVC End Point, the two SWVC End Points **MUST** be associated by the same SWVC.

9.3 SWVC List of UCSs Service Attribute

The value of the SWVC List of UCSs Service Attribute is a non-empty list of UCS Identifier Service Attribute values (section 12.1.1). The list contains one UCS Identifier value for each Underlay Connectivity Service that is used by the SD-WAN Service (SWVC).

- [R18] A UCS Identifier **MUST NOT** appear more than once in the value of the SWVC List of UCSs Service Attribute.

Without at least one entry in the value of the SWVC List of UCSs Service Attribute, an SWVC cannot transport any IP Packets and is therefore not useful.

9.4 SWVC Reserved Prefixes Service Attribute

The SWVC Reserved Prefixes Service Attribute specifies a list of IP Prefixes in the Subscriber's addressing domain that the Service Provider reserves for use for the SWVC within its own network or for distribution to the Subscriber via DHCP or SLAAC. The list can be empty or can contain IPv4 or IPv6 Prefixes or both. These IP Prefixes need to be agreed upon so as to ensure they do not overlap with IP Prefixes assigned by the Subscriber inside the Subscriber Network.

Note that it is not necessary to reserve the Service Provider's IP address on the directly connected subnet for a UNI using this attribute; such addresses are implicitly reserved via the SD-WAN UNI IPv4 Connection Addresses Service Attribute (section 11.4) and the SD-WAN UNI IPv6 Connection Addresses Service Attribute (section 11.5).

9.5 SWVC List of Zones Service Attribute

As described in section 7.12, the IP hosts in the Subscriber Network can be partitioned into Zones. Application Flows are defined at a UNI based on an Application Flow Specification and a Zone, so different Policies can be assigned to IP Packets that match the same Application Flow Specification but are in different Zones.

The SWVC List of Zones Service Attribute provides the mechanism to identify the Zones supported by the SD-WAN Service and to enumerate a set of Zone-wide Policies for each Zone. A Zone-wide Policy associates a Policy with an Application Flow Specification (or Application Flow Specification Group). At each UNI that has an SWVC End Point, the Zone-wide Policy is assigned to the Application Flow that matches the Zone and Application Flow Specification unless overridden by using the SWVC End Point Policy Map Service Attribute (section 10.4). For example, a Zone-wide Policy can be used to send all web traffic (e.g., ports 80 and 443) from the Guest-Wifi Zone to an Internet Breakout.

The Zone definition includes a list of IP prefixes that defines the IP subnets that compose the Zone. A host is assigned to a Zone if the IP address of the host is a member of one of the IP Prefixes for the Zone. An Ingress IP Packet is assigned to a Zone based on its source IP address. An IP Packet directed toward an Egress UNI is assigned a Zone based on its source IP address unless it arrives from the Internet (i.e., did not ingress at an SD-WAN UNI), in which case it is assigned the reserved Zone *Internet*.

The value of this Service Attribute is a non-empty list of 3-tuples $\langle zoneName, zonePrefixes, zoneIngressPolicies \rangle$ where:

- *zoneName* is an Identifier String that specifies the name of the Zone. It cannot have the values *Self*, *Internet*, or *All*.
- *zonePrefixes* is *Default* or a non-empty list of IP Prefixes that identifies the hosts in the Zone.
- *zoneIngressPolicies* is a list (possibly empty) of 2-tuples $\langle zoneIngressAFS, zoneIPol \rangle$ where:
 - *zoneIngressAFS* is an Application Flow Specification Name from the value of the SWVC List of Application Flow Specifications Service Attribute (see section 9.14) or an Application Flow Specification Group Name from the value of the SWVC List of Application Flow Specification Groups Service Attribute (see section 9.13).
 - *zoneIPol* is the name of an Ingress Policy (see section 9.12).

Note that the value of this Service Attribute is a non-empty list, so definition of at least one Zone is required in each SWVC. In this case, the Zone can map all IP Packets or a subset of them.

[R19] A Zone name, *zoneName*, **MUST** appear at most once in the value of the SWVC List of Zones Service Attribute.

The reserved Zone name *Internet* is assigned to IP Packets entering the SWVC directly from the Internet (i.e., not from a UNI). It allows Egress Policies to be assigned to Application Flows coming directly from an Internet UCS.

- [R20] An IP Prefix listed in *zonePrefixes* associated with a specific *zoneName* **MUST NOT** overlap with an IP Prefix listed in *zonePrefixes* associated with another *zoneName*.
- [R21] If the value of *zonePrefixes* in any 3-tuple is *Default*, then any IP host in the Subscriber Network whose IP address is not included in any *zonePrefix* list **MUST** be associated with the Zone whose *zonePrefix* element has the value *Default*.
- [R22] The value of the SWVC List of Zones Service Attribute **MUST NOT** contain more than one 3-tuple that has the value *Default* in the *zonePrefixes* element.
- [R23] If the IP host associated with the source IP address of an Ingress IP Packet is not assigned to a Zone, the IP Packet **MUST** be discarded.

Note that no IP Packets are discarded due to [R23] if one of the defined Zones has a *zonePrefix* with value *Default*.

- [R24] An IP Packet that arrives from a UNI **MUST** be associated with the Zone of the IP host associated with the Source IP Address of the IP Packet.
- [R25] An IP Packet that arrives from the Internet **MUST** be assigned to the reserved Zone *Internet*.
- [R26] An Application Flow Specification Name or Application Flow Specification Group Name, *zoneIngressAFS*, **MUST** appear at most once in the list of Zone-wide Policies, *zoneIngressPolicies* for each *zoneName* in the value of the SWVC List of Zones Service Attribute.

An example of the value of this Service Attribute might be:

```
[{"Guest-Wifi", [192.168.1.0/24, 192.168.2.0/24], [{"webflows", "Ibpolicy"}]},
{"Corp", [192,168.3.0/24], [{"webflows", "Ibpolicy"}, {"vnc", "Block"}, {"else", "private"}]}]
```

In this example, two Zones are defined, *Guest-Wifi* and *Corp*. The *Guest-Wifi* Zone has one Zone-wide Policy mapping, *Ibpolicy*, to forward all web traffic (*webflows*) directly to the Internet (via Internet Breakout). The *Corp* Zone has three global Policy mappings: all web traffic is forwarded to the Internet via *Ibpolicy*; all screen-sharing traffic using the *vnc* protocol is blocked via the *Block* Policy (see 9.12); and everything *else* is sent over a UCS that is a private network, via *private* Policy.

9.6 SWVC List of Virtual Topologies Service Attribute

As described in section 7.13, an SD-WAN Service can present different forwarding topologies to different Application Flows. The SWVC List of Virtual Topologies Service Attribute provides the mechanism for defining these different forwarding topologies. The value of this Service Attribute is a list of 3-tuples of the form $\langle vtName, vtType, vtEP \rangle$ where:

- *vtName* is an Identifier String that specifies the name of the Virtual Topology.

- *vtType* provides a description of the type of connectivity provided by the Virtual Topology. *vtType* can have one of the following values: *multipoint-to-multipoint*, *rooted-multipoint*.
- *vtEP* provides details of the SWVC End Points that are connected by the Virtual Topology. The form of *vtEP* is different for each value of *vtType*.

[R27] A Virtual Topology name, *vtName*, **MUST** appear exactly once in the value of the SWVC List of Virtual Topologies Service Attribute.

[R28] A Virtual Topology name, *vtName*, **MUST NOT** have the reserved name *Swvc*.

The reason for [R28] is that every SWVC has a pre-defined multipoint-to-multipoint Virtual Topology named *Swvc* that connects all SWVC End Points in the SWVC. Therefore, the value of this Service Attribute can be an empty list.

Specific details on each Virtual Topology type and its associated *vtEP* are described in sections 9.6.1 and 9.6.2. Section 9.12.2.4 includes requirements that define the specific forwarding constraints associated with each *vtType*.

9.6.1 *vtType=multipoint-to-multipoint*

The value *multipoint-to-multipoint* for *vtType* indicates that the Virtual Topology allows IP Packets in an Application Flow to be forwarded from any SWVC End Point specified in the *vtEP* element to any other SWVC End Point specified in the *vtEP* element.

[R29] If the value of *vtType* is *multipoint-to-multipoint* for a Virtual Topology in the SWVC List of Virtual Topologies Service Attribute, the value for *vtEP* **MUST** be a list of at least two SWVC End Point Identifiers (section 10.1).




Figure 14 – Example of two Multipoint-to-Multipoint Virtual Topologies

Figure 14 shows an example of two multipoint-to-multipoint Virtual Topologies in an SD-WAN network with four SD-WAN Edges. The red solid lines show a multipoint-to-multipoint Virtual Topology that contains all SWVC End Points, and the green dotted lines show another multipoint-to-multipoint Virtual Topology that contains End Points *Store 1*, *Store 3*, and *Data Center*.

9.6.2 *vtType=rooted-multipoint*

In the case of *multipoint-to-multipoint*, all the SWVC End Points included in the topology have full connectivity to each other. A *rooted-multipoint* Virtual Topology, on the other hand, has restricted connectivity, i.e., traffic can be forwarded between some pairs of SWVC End Points, but not other pairs. This topology is often called a *tree* topology. Using this latter analogy, a rooted-multipoint Virtual Topology has some End Points referred to as *root* SWVC End Points and some SWVC End Points referred to as *leaf* SWVC End Points.

Root SWVC End Points have no forwarding restrictions, i.e., they can communicate with all other root SWVC End Points as well as all leaf SWVC End Points in the Virtual Topology. Leaf SWVC End Points, on the other hand, can only communicate with root SWVC End Points.

- [R30] If the value of *vtType* is *rooted-multipoint* for a Virtual Topology in the SWVC List of Virtual Topologies Service Attribute, the value of *vtEP* for the Virtual Topology **MUST** be a 2-tuple $\langle \text{rootlist}, \text{leaflist} \rangle$, where *rootlist* is a non-empty list of SWVC End Point Identifiers of root SWVC End Points and *leaflist* is a non-empty list of SWVC End Point Identifiers of leaf SWVC End Points (section 10.1).
- [R31] If the value of the *vtType* is *rooted-multipoint* for a Virtual Topology in the SWVC List of Virtual Topologies Service Attribute, an SWVC End Point Identifier **MUST** appear at most once in either the list of root SWVC End Points or the list of leaf SWVC End Points.

As noted in [R30], there must be at least one root SWVC End Point and one leaf SWVC End Point in the Virtual Topology. A tree with just two SWVC End Points is not very interesting (it is, in effect, just a multipoint-to-multipoint topology with two End Points), but additional SWVC End Points can be added.




Figure 15 – Example of Rooted-multipoint Virtual Topology

Figure 15 shows an example of a *rooted-multipoint* Virtual Topology in an SD-WAN Service with four SD-WAN Edges. There is only a single root SWVC End Point, but there could be more than one. For example, there could be a backup POS controller at another location. The two controllers could communicate with each other since they would both be at root SWVC End Points, and all three stores (leaf SWVC End Points) could communicate with both controllers.

9.7 SWVC Performance Time Intervals Service Attribute

The SWVC Performance Time Intervals Service Attribute specifies a set of times and time intervals used in the computation of Performance Metrics or information rates.

The value of the Service Attribute is a 3-tuple $\langle evalinterval, arrivalinterval, irduration \rangle$ where:

- *evalinterval* is the interval in milliseconds over which the Performance Metrics specified in the PERFORMANCE Policy Criterion (section 9.12.2.8) are evaluated.
- *arrivalinterval* is the difference in arrival times (specified in milliseconds) at the Ingress UNI between two packets used to compute the Mean One-Way Packet Delay Variation (as described in section 13.4).
- *irduration* is the time interval in milliseconds over which the information rate is determined in the evaluation of the BANDWIDTH Policy Criterion (section 9.12.2.10). The information rate is determined over any time interval of duration *irduration*—i.e., it uses a ‘sliding window’ rather than a fixed, recurring, window.

[D2] The *irduration* used for the BANDWIDTH Policy Criterion **SHOULD NOT** exceed 1000 milliseconds.

9.8 SWVC Service Performance Objectives Reporting Periods Service Attribute

Service Performance Objectives (see sections 11.7 and 11.8) are evaluated over a series of consecutive time periods. The value of this Service Attribute is a 2-tuple $\langle s, T \rangle$ where:

- *s* is a time that represents the date and time that evaluation of Service Performance Objectives starts for the SWVC and its components.
- *T* is a time duration, e.g., 1 month or 2 weeks, that is used in conjunction with *s* to specify time intervals for determining when Service Performance Objectives are met.

The time periods are specified by elements *s* and *T*. One time period, denoted T_0 , starts at time *s* and has duration *T*. Each subsequent time period, denoted T_k , starts at time $s + kT$ where *k* is an integer, and has duration *T*; in other words, each new time period starts as soon as the previous one ends. Each Service Performance Objective is evaluated for each time period T_k , so one can say that for a given T_k , the Service Performance Objective is either met or not met.

Note that *T* can be specified using any time units; in particular, calendar months are allowable. In this case, if *s* is specified as, for example, midnight on January 5, 2021, and *T* is 1 calendar month, then each subsequent T_k will start at midnight on the 5th of the month. The details of possible values of *s* and *T* are beyond the scope of this document but, like all Service Attribute values, need to be agreed to by the Subscriber and Service Provider.

9.9 SWVC List of Rate Limiters Service Attribute

The SWVC List of Rate Limiters Service Attribute specifies the named Rate Limiters (see section 7.15) that can be referenced in the BANDWIDTH Policy Criterion (see section 9.12.2.10).

Each named Rate Limiter is instantiated, at most, once at an SWVC End Point and the set of Application Flows whose Policy refers to a named Rate Limiter (referred to as a Bandwidth Flow) share the bandwidth constraints specified by the parameters of the Rate Limiter.

The value of the SWVC List of Rate Limiters Service Attribute is a list (which may be empty) of 3-tuples, $\langle name, commit, limit \rangle$ where:

- *name* is an Identifier String that is the name of the Rate Limiter.
- *commit* is the threshold information rate (bits per second) at or below which the SD-WAN Service Provider commits to deliver packets in the Bandwidth Flow with high probability under all traffic conditions.
- *limit* is the threshold information rate (bits per second) above which the Service Provider does not deliver IP Packet in the Bandwidth Flow under any traffic conditions.

Following is an example of the value of this Service Attribute:

```
[  
<RL1, 100Mbps, 200Mbps>  
<RL2, 50Mbps, 250Mbps>  
]
```

9.10 SWVC List of Security Policies Service Attribute

The value of the SWVC List of Security Policies Service Attribute is a list (which may be empty) of Security Policies that are available for use with this SWVC. MEF 88 [38] section 6 defines the structure and content of each Security Policy in the SWVC List of Security Policies Service Attribute. The Security Policies are assigned to Ingress and Egress Application Flows via the AF-SECURITY-INGRESS Policy Criterion (section 9.12.2.11) and the AF-SECURITY-EGRESS Policy Criterion (section 9.12.3.2). One of the Security Policy parameters defined in MEF 88 [38] is the Security Policy Identifier.

9.11 SWVC List of Application Flow Business Importance Levels Service Attribute

Applications, and hence their Application Flows, have different levels of business importance to the Subscriber. The SWVC List of Application Flow Business Importance Levels Service Attribute specifies an ordered list of labels that can be assigned to Application Flows to indicate the relative business importance of each Application Flow.

The value of this Service Attribute is *None* or a 2-tuple of the form $\langle implist, default \rangle$ where:

- *implist* is a non-empty list of Identifier Strings that are the business importance labels ordered from highest business importance to lowest.
- *default* is one of the labels in *implist* that represents the business importance of Application Flows that do not have a business importance assigned by Policy.

In general, it is expected that the Subscriber and Service Provider have agreed on sufficient resources in the SD-WAN Service (e.g., number of UCSs, UCS bandwidth, number of Paths between SD-WAN Edges, etc.) to meet the service commitments indicated by the Policies assigned to Application Flows. In spite of this, there might be situations where these commitments cannot be met — large simultaneous bursts of ingress traffic, network component failures, re-routing, etc.

In these cases, the SD-WAN Service can experience resource contention that might result in the need to delay or discard packets. During these periods, the Service Provider might not be able to meet the SD-WAN Service Policy commitments. Rather than allowing random decisions about which packets to delay or discard, each Policy specifies (explicitly or implicitly) a relative business importance of the Application Flow to the Subscriber via the BUSINESS-IMPORTANCE Policy Criterion (see section 9.12.2.8) and thereby provides guidance to the Service Provider in these contention situations.

Examples of the value of this Service Attribute are:

- $\langle [best, ok, worst], ok \rangle$ - indicates three levels of business importance and “ok” is the *default* level of importance.
- $\langle [import], import \rangle$ - indicates one level of business importance.

The value *None* and a list with one element have the same effect, a single level of business importance for all Application Flows. However, if the value is *None*, Policies in the SD-WAN Service cannot include the BUSINESS-IMPORTANCE Policy Criterion (see [R66]), but if the list includes a single element, the BUSINESS-IMPORTANCE Policy Criterion can be included in Policies, thus providing a basis for future expansion of the number of levels.

9.12 SWVC List of Policies Service Attribute

Associated with each SWVC is a list of named Policies that can be assigned to each Application Flow using the SWVC End Point Policy Map (see section 10.4), or via Zone-wide Policies specified in the SWVC List of Zones (see section 9.5).

A Policy is composed of Policy Criteria. Each Policy Criterion is either an Ingress Policy Criterion—meaning that it is applicable to Ingress Application Flows—or an Egress Policy Criterion—meaning that it is applicable to Egress Application Flows—and each Policy can only contain one type (Ingress or Egress). If a Policy includes Ingress Policy Criteria, it is referred to as an Ingress Policy, and if a Policy includes Egress Policy Criteria, it is referred to as an Egress Policy.

The value of this Service Attribute is a non-empty list of 2-tuples of the form $\langle polName, polCL \rangle$ where:

- *polName* is an Identifier String that specifies the name of the Policy. *polName* cannot be “None” or the reserved Policy name “Block”.
- *polCL* is a non-empty list of Policy Criteria 2-tuples, of the form $\langle Pcname, Pcparam \rangle$ where:
 - *Pcname* is an Identifier String containing a Policy Criterion name from Table 5 or Table 6, or a Service Provider-defined Policy Criterion name.
 - *Pcparam* is a parameter value specific to the Policy Criterion specified in *Pcname*.

[R32] A Policy name, *polName*, in the value of the SWVC List of Policies Service Attribute **MUST** appear, at most, once.

The reserved *polName*, *Block*, must not appear in the value of the SWVC List of Policies Service Attribute. *Block* is a pre-defined Ingress Policy that results in the discard of all IP Packets in any Application Flow to which it is assigned.

- [R33] A Policy Criterion name, *Pcname*, **MUST** appear, at most, once in each list of Policy Criteria, *polCL*, in the value of the SWVC List of Policies Service Attribute.
- [R34] A Policy **MUST** contain either Ingress Policy Criteria or Egress Policy Criteria but not both.
- [R35] The Policy Criteria supported by the Service Provider for composing SD-WAN Policies **MUST** include the Policy Criteria listed in Table 5.

Pcname	Ingress / Egress	Description	Values
ENCRYPTION	Ingress	Indicates whether the Application Flow requires encryption	<i>Required-Always, Required-Public-Only, or Optional</i>
INTERNET-BREAKOUT	Ingress	Indicates whether IP Packets in the Application Flow can be forwarded to Internet destinations	<i>Enabled, Disabled</i>
PUBLIC-PRIVATE	Ingress	Indicates whether the Application Flow can traverse Internet Access Underlay Connectivity Services or not	<i>Private-only, Either</i>

Table 5 – Policy Criteria – Support Required

- [D3] The Policy Criteria supported by the Service Provider for composing SD-WAN Policies **SHOULD** include the Policy Criteria listed in Table 6.

Pcname	Ingress / Egress	Description	Values
BACKUP	Ingress	Indicates whether this Application Flow can use a UCS designated as “backup”	<i>Enabled, Disabled</i>
VIRTUAL-TOPOLOGY	Ingress	Indicates the Virtual Topology that the Application Flow should be forwarded over	Virtual Topology Name
ALLOWED-DESTINATION-ZONES	Ingress	Specifies which Zones the IP Packets in the Application Flow can be delivered to	<i>All</i> or a list of Zone Names
BILLING-METHOD	Ingress	Indicates whether the Application Flow can be sent over an Underlay Connectivity Service that has usage-based or flat-rate billing	<i>Flat-Rate-Only, Usage-Based-Only, Either</i>
BUSINESS-IMPORTANCE	Ingress	Indicates the relative business importance of the Application Flow	A business importance label
PERFORMANCE	Ingress	Specifies a list of performance requirements for the Application Flow	See section 9.12.2.8
BANDWIDTH	Ingress	Specifies a bandwidth commitment and bandwidth limit on the Application Flow.	A Rate Limiter
AF-SECURITY-INGRESS	Ingress	Specifies the Security Functions to apply to the Application Flow at the Ingress UNI	<i>None</i> or an Application Flow Security Policy Identifier
EGRESS-BLOCK	Egress	Indicates whether the Application Flow is allowed to be forwarded towards the Egress UNI	<i>Allow or Discard</i>
AF-SECURITY-EGRESS	Egress	Specifies the Security Functions to apply to the Application Flow at the Egress UNI	<i>None</i> or an Application Flow Security Policy Identifier

Table 6 – Policy Criteria – Support Recommended

Note that the Policy Criteria names used in Table 5 and Table 6 and the syntax used in the examples below are intended to describe the behavior of each Policy Criterion, and not to mandate an implementation or syntax. For example, no implementation needs to have a command or configuration parameter called BILLING-METHOD, but rather it can provide Policy functionality consistent with the BILLING-METHOD Policy Criterion described in this document.

The behavior of these Policy Criteria is described in subsequent sections.

- [R36]** If the Service Provider defines its own Policy Criteria, the description of each Policy Criterion agreed upon with the Subscriber **MUST** include the following items:
- A name for the Policy Criterion
 - The possible values for the Policy Criterion
 - The behavior associated with each value
 - Whether it is an Ingress Policy Criterion or an Egress Policy Criterion
 - The behavior of each value when used with INTERNET-BREAKOUT (for Ingress Policy Criteria)

- Any interactions that the Policy Criterion has with other Policy Criteria

[R37] If the Service Provider defines its own Policy Criterion as described in [R36], it **MUST NOT** reuse the *Pcnames* in Table 5 and Table 6.

Example Ingress Policies (ipA and ipB) and Egress Policies (epA and epB) are shown below:

```
[
(ipA, [ <ENCRYPTION, Required-Public-Only >,
        <INTERNET-BREAKOUT, Disabled>,
        <PUBLIC-PRIVATE, Either>,
        <BILLING-METHOD, Flat-rate-only>,
        <BACKUP, Disabled>,
        <BUSINESS-IMPORTANCE, Super>,
        <VIRTUAL-TOPOLOGY, "Swvc" >,
        <BANDWIDTH, "Rlimit1" >
]),
```

```
(ipB, [ <ENCRYPTION, Required-Always>,
        <INTERNET-BREAKOUT, Disabled>,
        <PUBLIC-PRIVATE, Private-only>,
        <BILLING-METHOD, Flat-rate-only>,
        <BACKUP, Enabled>,
        <VIRTUAL-TOPOLOGY, "EngFacilities" >,
        <BANDWIDTH, EngBW>
        <AF-SECURITY-INGRESS "Paranoid">,
]),
```

```
<epA, [(EGRESS-BLOCK, "Allow")]>],
```

```
<epB, [(EGRESS-BLOCK, "Discard")]>
]
```

9.12.1 Policy Criteria specification and interaction

Each Service Provider supports a set of Policy Criteria that can include both criteria from the lists in Table 5 and Table 6 as well as other criteria defined by the Service Provider (see [R36]). For a given SD-WAN Service, the Subscriber and the Service Provider agree (via this Service Attribute) on the criteria that will be used. This may be the entire set of Policy Criteria supported by the Service Provider or a subset.

[R38] For an Ingress IP Packet in a given Application Flow, if the Service Provider cannot forward the packet to the destination over an Underlay Connectivity Service (or a sequence of Underlay Connectivity Services) that meets the Policy for that Application Flow, the packet **MUST** be discarded.

Requirement [R38] indicates two necessary conditions that must be met for an IP Packet to be forwarded by the SD-WAN Service:

- A forwarding path exists, and
- The forwarding path meets the Policy assigned to the Application Flow

A few of the Policy Criteria descriptions in the following section refer to a TVC or UCS as being available or not available. Being “available” in this context means that it meets the two [R38] conditions. It is the Service Provider’s responsibility to ensure that these conditions are met, but there can be transient and failure situations when they are not met.

Many of the Policy Criteria (but not all) are used to specify conditions that restrict the TVCs that are eligible to carry the Application Flow. Most of these Policy Criteria are binary, i.e., a TVC (in the Path used to forward the Application Flow) is eligible or not. The intersection of TVCs that satisfy these Policy Criteria provides an initial determination of the Paths that can carry the Application Flow.

There are, however, some Policy Criteria that are not binary. For example, the PERFORMANCE Policy Criterion results in an ordered list of Paths based on a determination of how well they meet the stated performance goals for the Application Flow. The BANDWIDTH Policy Criterion is one of the few Policy Criteria that are more focused on individual IP Packets in the Application Flow rather than the Application Flow itself.

9.12.2 Ingress Policy Criteria

This section describes the parameters and behavior of the Ingress Policy Criteria specified in this document.

9.12.2.1 ENCRYPTION Ingress Policy Criterion

The ENCRYPTION Policy Criterion provides a mechanism to specify whether an Application Flow is required to traverse encrypted TVCs. It can have the value *Required-Always*, *Required-Public-Only*, or *Optional*.

- [R39] If Policy Criterion ENCRYPTION=*Required-Always* is applied to an Application Flow, the Path chosen for forwarding the Application Flow **MUST** include only TVCs that encrypt IP Packets in the Application Flow before forwarding over the Underlay Connectivity Service.
- [R40] If Policy Criterion ENCRYPTION=*Required-Public-Only* is applied to an Application Flow, any TVCs in the Path chosen for forwarding the Application Flow that traverse Public UCSs **MUST** encrypt IP Packets in the Application Flow before forwarding over the Public UCS.

Requirements [R39] and [R40] mean that encryption is applied by the SD-WAN Edge before packets are forwarded over the UCS UNI, and decryption is applied to packets received from the UCS at the destination SD-WAN Edge.⁹

- [R41] If an Ingress Policy that does not include the ENCRYPTION Policy Criterion or an Ingress Policy that includes the ENCRYPTION Policy Criterion with the value *Optional* is assigned to an Application Flow, then ENCRYPTION

⁹ Requirements related to the level and type of encryption are out of scope but may be addressed in a future version of the specification.

MUST NOT be considered in the forwarding decision for the Application Flow.

Optional indicates that the Subscriber doesn't need encryption for the Application Flow, and therefore it can be sent over either encrypted or unencrypted TVCs at the Service Provider's discretion.

9.12.2.2 PUBLIC-PRIVATE Ingress Policy Criterion

An SD-WAN Service can use *private* Underlay Connectivity Services such as MEF Carrier Ethernet Services and *public* Underlay Connectivity Services, i.e., Internet Access Services. The PUBLIC-PRIVATE Policy Criterion provides control over whether or not an Application Flow can traverse a public Underlay Connectivity Service, i.e., the Internet. It can have the value *Private-Only* or *Either*.

- [R42] If Policy Criterion PUBLIC-PRIVATE=*Private-Only* is applied to an Application Flow, the Path chosen for forwarding the Application Flow **MUST** include only TVCs that traverse a UCS whose UCS Type Service Attribute=*Private* (section 12.1.2).
- [R43] If an Ingress Policy that does not include the PUBLIC-PRIVATE Policy Criterion or an Ingress Policy that includes the PUBLIC-PRIVATE Policy Criterion with the value *Either* is assigned to an Application Flow, then PUBLIC-PRIVATE **MUST NOT** be considered in the forwarding decision for the Application Flow.

9.12.2.3 BILLING-METHOD Ingress Policy Criterion

The cost for the use of a particular Underlay Connectivity Service can be flat rate (i.e., based on units of time such as €500/month) or usage-based (i.e., based on how much data is sent across it such as €10/TB). The BILLING-METHOD Policy Criterion provides control over the charge type of the network that can be used to forward an Application Flow. It can have the value *Flat-Rate-Only*, *Usage-Based-Only*, or *Either*.

- [R44] If Policy Criterion BILLING-METHOD=*Flat-Rate-Only* is applied to an Application Flow, the Path chosen for forwarding the Application Flow **MUST** traverse only UCSs whose UCS Billing Method Service Attribute=*Flat-rate* (section 12.1.3).
- [R45] If Policy Criterion BILLING-METHOD=*Usage-Based-Only* is applied to an Application Flow, the Path chosen for forwarding the Application Flow **MUST** traverse only UCSs whose UCS Billing Method Service Attribute=*Usage-based* (section 12.1.3).
- [R46] If an Ingress Policy that does not include the BILLING-METHOD Policy Criterion or an Ingress Policy that includes the BILLING-METHOD Policy Criterion with the value *Either* is assigned to an Application Flow, then

BILLING-METHOD **MUST NOT** be considered in the forwarding decision for the Application Flow.

Refer to MEF 74 [37] for examples of a broader range of billing options that are currently beyond the scope of this document.

9.12.2.4 VIRTUAL-TOPOLOGY Ingress Policy Criterion

The value of the VIRTUAL-TOPOLOGY Ingress Policy Criterion is either a Virtual Topology name from the *vtName* element in the value of the SWVC List of Virtual Topologies Service Attribute (section 9.6) or the value *Swvc*. The value specifies the Virtual Topology used to constrain forwarding of Ingress Application Flows to which it is applied.

Each Virtual Topology includes a list (or possibly two lists) of SWVC End Point Identifiers (see section 9.6), and the SWVC End Point is “in” the Virtual Topology if its Identifier is included in one of these lists.

- [R47] A Policy **MUST NOT** be assigned to an Application Flow at a UNI if the SWVC End Point at that UNI is not in the Virtual Topology specified in the VIRTUAL-TOPOLOGY Policy Criterion in the Policy.

[R47] specifies the restriction on the SWVC End Point where a Virtual Topology is assigned to an Ingress Application Flow. Restrictions on the SWVC End Points at which the Application Flow can egress are specified independently for each of the Virtual Topology types. Note that [R47] applies regardless of how the Policy is assigned to the Application Flow (see section 10.4.1). One implication of [R47] is that if a Zone-wide Policy includes the VIRTUAL-TOPOLOGY Policy Criterion, the Zone-wide Policy must be overridden at SWVC End Points that are not included in the specified Virtual Topology.

- [R48] If the Ingress Policy assigned to an Application Flow does not include the VIRTUAL-TOPOLOGY Policy Criterion, the effect of the Policy **MUST** be as if *VIRTUAL-TOPOLOGY=Swvc* were included in it.
- [R49] When a Policy containing the VIRTUAL-TOPOLOGY Policy Criterion is assigned to an Application Flow at an SWVC End Point and the value of *vtType* that is associated with the specified *vtName* is *multipoint-to-multipoint*, IP Packets in that Application Flow **MUST** egress only at SWVC End Points that are in the list of SWVC End Points for the Virtual Topology.

The *Swvc* Virtual Topology (see [R48]) is a multipoint-to-multipoint Virtual Topology that includes all SWVC End Points, and therefore an Ingress IP Packet assigned to the *Swvc* Virtual Topology can egress at any other SWVC End Point in the SWVC.

- [R50] When a Policy containing the VIRTUAL-TOPOLOGY Policy Criterion is assigned to an Application Flow at an SWVC End Point:
- if the value of *vtType* that is associated with the specified *vtName* is *rooted-multipoint*, and

- if the SWVC End Point is in the list of leaf End Points in the Virtual Topology,

Then IP Packets in that Application Flow **MUST** egress only at SWVC End Points that are in the list of root End Points for the Virtual Topology.

- [R51]** When a Policy containing the VIRTUAL-TOPOLOGY Policy Criterion is assigned to an Application Flow at an SWVC End Point:
- if the value of *vtType* that is associated with the specified *vtName* is *rooted-multipoint*, and
 - if the SWVC End Point is in the list of root End Points in the Virtual Topology,

Then IP Packets in that Application Flow **MUST** egress only at SWVC End Points that are in the list of root SWVC End Points or the list of leaf SWVC End Points for the Virtual Topology.

Requirements [R50] and [R51] describe the forwarding restrictions associated with a rooted-multipoint Virtual Topology. [R50] specifies that an IP Packet that arrives from the Subscriber Network at a leaf SWVC End Point can only egress at a root SWVC End Point that is in the Virtual Topology, whereas [R51] indicates that if the ingress is at a root SWVC End Point, it can egress at any End Point in the Virtual Topology, root or leaf.

9.12.2.5 INTERNET-BREAKOUT Ingress Policy Criterion

The INTERNET-BREAKOUT Policy Criterion indicates whether IP Packets in the Application Flow can be forwarded to an Internet destination by the SD-WAN Service (see section 7.14). It can have the value *Enabled* or *Disabled*. If an Application Flow is assigned a Policy that includes INTERNET-BREAKOUT=*Enabled*, then each IP Packet in the Application Flow can be forwarded to the Internet or to another UNI in the SWVC based on routing information and its destination address. Note that forwarding to another UNI can be precluded by including the ALLOWED-DESTINATION-ZONES Policy Criterion with an empty list (see section 9.12.2.6).

Terminology Note: In the following requirements the term “Internet destination” is used to refer to an IP host whose IP address is reachable on the Internet via an Internet Access UCS rather than via a UNI in the SWVC. An “Internet destination” should not be confused with an IP host that is reachable via an SD-WAN UNI over a TVC that traverses an Internet UCS.

- [R52]** If the Policy Criterion INTERNET-BREAKOUT=*Enabled* is applied to an Application Flow, IP Packets in the Application Flow that are addressed to Internet destinations **MUST** be forwarded to the Internet over an Internet Access UCS.

Consistent with [R178], an IP Packet that is forwarded to the Internet via an Internet Access UCS must be forwarded at a UCS End Point with the UCS End Point Breakout Service Attribute = *Enabled* (see section 12.3.3).

The Service Provider chooses an appropriate Internet Access UCS over which to forward Internet-destined IP Packets in the Application Flow. This may be an Internet Access UCS connected

directly to the SD-WAN Edge where the Ingress UNI for the Application Flow is located (referred to as Local Internet Access), or it may be an Internet Access UCS located at another SD-WAN Edge (in which case the IP Packets are first sent over one or more TVCs to that SD-WAN Edge).

- [R53] If the Policy Criterion INTERNET-BREAKOUT=*Disabled* is applied to an Application Flow, IP Packets in the Application Flow that are addressed to Internet destinations **MUST** be discarded.
- [R54] If the Policy Criterion INTERNET-BREAKOUT=*Enabled* is applied to an Application Flow, all the other Policy Criteria specified in Table 5 and Table 6 except BANDWIDTH, BILLING-METHOD, and AF-SECURITY-INGRESS, **MUST** be ignored for IP Packets in the Application Flow that are forwarded to the Internet via Local Internet Breakout.

Requirement [R54] acknowledges that Policy Criteria other than those specified in the requirement (e.g., PUBLIC-PRIVATE and ENCRYPTION) are not relevant for Local Internet Breakout. Application Flows for packets destined for the Internet can, however, have Bandwidth limitations. If the Internet Breakout occurs at an SD-WAN Edge other than the local one, the other Policy Criteria are used to describe forwarding of IP Packets within the SD-WAN Service until the breakout location is reached.

- [R55] If the Policy Criterion INTERNET-BREAKOUT=*Enabled* is applied to an Application Flow, the VIRTUAL-TOPOLOGY Policy Criterion **MUST** be ignored for IP Packets in the Application Flow that are addressed to Internet destinations.

Requirement [R55] acknowledges that since the Internet is not an SWVC End Point, it is not part of any Virtual Topology and therefore, the VIRTUAL-TOPOLOGY Policy Criterion is not relevant for forwarding packets that are subject to Internet Breakout.

- [R56] If no Application Flows at a UNI are assigned a Policy that includes INTERNET-BREAKOUT=*Enabled*, then IP Packets from the Internet that are destined to an IP host via the UNI **MUST** be discarded.

The purpose of [R56] is to protect the Subscriber Network from unsolicited Internet traffic. If none of the Application Flows at the UNI can “break out” to the Internet, then no traffic *from* the Internet should be forwarded back to the Subscriber at the UNI.

- [R57] If the Ingress Policy assigned to an Application Flow does not include the INTERNET-BREAKOUT Policy Criterion, the effect of the Policy **MUST** be as if the Policy includes the INTERNET-BREAKOUT Policy Criterion with value *Disabled*.

9.12.2.6 ALLOWED-DESTINATION-ZONES Ingress Policy Criterion

The ALLOWED-DESTINATION-ZONES Ingress Policy Criterion indicates the Zones that an Ingress Packet can be delivered to, i.e., the Zones of destination IP hosts as determined by the

destination IP addresses. The value is either *All* or a list where each entry is either the reserved name *Self* or a Zone name (see section 9.5).

The reserved name *Self* is equivalent to the Zone that the IP Packet is in (i.e., the Zone of the source IP host). This allows a Policy to be reused for different Zones. Hence, ALLOWED-DESTINATION-ZONES [*Self*] means that the IP Packet can only be delivered to an Egress UNI if the destination address in the IP Packet is in the same Zone as the source address. Similarly, ALLOWED-DESTINATION-ZONES [*corp*, *Self*] means that the IP Packet can only be delivered to an Egress UNI if the destination address in the IP Packet is in the same Zone as the source address or is in Zone “corp”. Note that, in some cases, *Self* could actually be “corp”, but this doesn’t have any impact on the Policy Criterion.

- [R58]** If the value of the ALLOWED-DESTINATION-ZONES Policy Criterion is a list of Zone names, each entry in the list **MUST** either be *Self* or be in the value of the SWVC List of Zones Service Attribute (section 9.5).

Note that requirement [R58] implies that the reserved Zone name *Internet* cannot be used as a value for this Policy Criterion since this Zone is used only for IP Packets arriving from the Internet, to which no Ingress Policy is applied.

- [R59]** If the value of the INTERNET-BREAKOUT Policy Criterion (see section 9.12.2.5) is *Disabled*, the value of the ALLOWED-DESTINATION-ZONES Policy Criterion **MUST NOT** be an empty list.

If Internet Breakout is *Enabled* for the Application Flow, an empty list for ALLOWED-DESTINATION-ZONES can be used to ensure that IP Packets in the Application Flow are delivered only to Internet destinations and not to other UNIs in the SWVC. But if Internet Breakout is *Disabled* for the Application Flow, then an empty list would result in all IP Packets in the Application Flow being discarded.

- [R60]** If the value of the ALLOWED-DESTINATION-ZONES Policy Criterion is a list of Zone names, each entry in the list **MUST** be unique.

Requirement [R60] requires that each Zone name (including *Self*) can appear in the list exactly once.

- [R61]** A unicast IP Packet in an Application Flow **MUST NOT** be forwarded to an Egress UNI unless the value of the ALLOWED-DESTINATION-ZONES Policy Criterion is *All*, or the Destination Address in the IP Packet maps to either:
- one of the Zones listed in the ALLOWED-DESTINATION-ZONES Ingress Policy Criterion; or
 - the Zone of the source IP host, if the ALLOWED-DESTINATION-ZONES Policy Criterion includes *Self*.

Requirement [R61] does not preclude IP Packets from being forwarded to Internet destinations if INTERNET-BREAKOUT=*Enabled*, since such IP Packets are not forwarded over a UNI.

- [R62] If the Ingress Policy assigned to an Application Flow does not include the ALLOWED-DESTINATION-ZONES Policy Criterion, the effect of the Policy **MUST** be as if the Policy includes the ALLOWED-DESTINATION-ZONES Policy Criterion with value [*Self*].

Requirement [R62] indicates that if the Ingress Policy assigned to an Application Flow does not include the ALLOWED-DESTINATION-ZONES Policy Criterion, then each IP Packet in the Application Flow can be forwarded only to Internet destinations (if INTERNET-BREAKOUT=*Enabled*) or destinations that are in the same Zone as the IP Packet.

9.12.2.7 **BACKUP Ingress Policy Criterion**

A UCS can be designated as *Backup* at one or more of its UCS End Points (section 12.3.2). When there is at least one non-*Backup* UCS available at an SD-WAN Edge (availability of a UCS is defined by the explanatory text after [R38]), Application Flows are not forwarded toward the egress UNI over a UCS that is designated as *Backup*. However, since a *Backup* UCS may have lower bandwidth and/or higher cost, it may be desirable to restrict which Application Flows are permitted to use them. This control can be achieved using the BACKUP Policy Criterion. It can have the value *Enabled* or *Disabled*.

- [R63] Application Flows **MUST NOT** be forwarded toward the destination UNI over a UCS End Point where the UCS End Point Backup Service Attribute=*Enabled*, if a Path to the destination egress UNI is available over a non-*Backup* UCS (i.e., where that Service Attribute value is *Disabled*).
- [R64] If the Policy Criterion BACKUP=*Disabled* is assigned to an Application Flow, then IP Packets in the Application Flow **MUST** be discarded if only Paths to the destination egress UNI over *Backup* UCSs are available.
- [R65] If the Ingress Policy assigned to an Application Flow does not include the BACKUP Policy Criterion, the effect of the Policy **MUST** be as if the Policy includes the BACKUP Policy Criterion with value *Disabled*.

Note that Application Flows that have BACKUP=*Enabled* will likely be more resilient than those with BACKUP=*Disabled*.

9.12.2.8 **BUSINESS-IMPORTANCE Ingress Policy Criterion**

The BUSINESS-IMPORTANCE Ingress Policy Criterion indicates the relative business importance of this Application Flow compared with other Application Flows in the SD-WAN Service. This determination is relevant during periods of resource contention when the Service Provider might have to choose to delay or discard IP Packets. During these periods, the Service Provider should provide better service to Application Flows that have a higher business importance.

- [D4] During periods of resource contention, if it is necessary for the Service Provider to discard IP Packets or degrade their performance, these actions **SHOULD** target IP Packets in Application Flows that have a lower business importance.

The value of the BUSINESS-IMPORTANCE Policy Criterion is the name of a business importance level specified in the *implist* element in the value of the SWVC List of Application Flow Business Importance Levels Service Attribute (see section 9.11).

- [R66] If the value of the SWVC List of Application Flow Business Importance Levels Service Attribute is *None*, the BUSINESS-IMPORTANCE Policy Criterion **MUST NOT** be included in any Policy and all Application Flows are assumed to have the same level of business importance.
- [R67] If the value of the SWVC List of Application Flow Business Importance Levels Service Attribute is not *None* and the Ingress Policy assigned to an Application Flow does not include the BUSINESS-IMPORTANCE Policy Criterion, the effect of the Policy **MUST** be as if the Policy includes the BUSINESS-IMPORTANCE Policy Criterion with value equal to the *default* element in the value of the SWVC List of Application Flow Business Importance Levels Service Attribute.

9.12.2.9 PERFORMANCE Ingress Policy Criterion

One of the benefits of SD-WAN is that, assuming that there are multiple ways of reaching a destination, the SD-WAN Service can dynamically choose a Path that best meets the Policy applied to an Application Flow, and this includes Policy Criteria associated with performance.

An SD-WAN service can monitor the performance of the various Paths between SD-WAN Edges in real time and adjust the forwarding decisions based on the most recently measured performance. This document does not specify how or when an SD-WAN implementation measures performance or even that it does measure it; however, there is an expectation that the Service Provider can implement the PERFORMANCE Policy Criterion if it is included in a Policy.

The PERFORMANCE Policy Criterion allows the Subscriber to indicate the important Performance Metrics for each Ingress Application Flow. The value of the PERFORMANCE Policy Criterion is a 2-tuple $\langle primary, secondary \rangle$ or *None*, where:

- *primary* is a 4-tuple $\langle metric, threshold, ceiling, remediation \rangle$ that describes the most important Performance Metric for this Application Flow. *metric* is one of the Performance Metrics listed in Table 7 and *threshold*, *ceiling*, and *remediation* are values for this Performance Metric, each of which can be *None*.
- *secondary* is a list of zero or more 2-tuples $\langle metric, threshold \rangle$. *metric* is one of the Performance Metrics listed in Table 7 and *threshold* is a value for this Performance Metric described below.

- [R68] A Performance Metric *metric* **MUST** appear at most once in the value of the PERFORMANCE Policy Criterion.

[R68] indicates that each performance metric (e.g., One-Way Mean Packet Delay) can only be specified once in the value of the PERFORMANCE Policy Criterion, either in the *primary* element or one of the entries in the list of *secondary* elements.

The result of most Policy Criteria are binary decisions, i.e., either a Path meets the requirement or not, whereas the result of the PERFORMANCE Policy Criterion is described as an ordered list of acceptable Paths. The Performance Metrics specified in the value of the PERFORMANCE Policy Criterion are evaluated for each Path to the Egress UNI and the Paths are ordered by how well they meet the stated performance goals.

Performance Metric Name	Description
One-Way Mean Packet Delay	See section 13.3
One-Way Mean Packet Delay Variation	See section 13.4
One-Way Packet Loss Ratio	See section 13.5

Table 7 – Performance Metrics

Note: For the three Performance Metrics used in this document (listed in Table 7), higher values indicate worse performance, so the normative descriptions in this document use “greater than” to express worse performance and “less than” to express better performance.

The *threshold* element in both *primary* and *secondary* is a value for the Performance Metric that indicates the point at which the use of a Path for this Application Flow is not preferred, i.e., if the measured value of this Performance Metric is worse than the *threshold*, selection of the Path for this Application Flow should be avoided. This is indicated in [R69], [R70], and [D5].

[R69] When the value of the PERFORMANCE Policy Criterion is not *None*, a Path whose performance for any of the specified Performance Metrics (primary or secondary) is worse than the corresponding *threshold* value, if not *None*, **MUST** be placed in the list of Paths that result from the evaluation of this Policy Criterion after all Paths for which this is not the case.

A Path as described in [R69] is not necessarily put at the end of the list but must be put after all of the Paths whose performance is better than or equal to the threshold values (if not *None*) for all of the specified Performance Metrics.

[R70] When the value of the PERFORMANCE Policy Criterion is not *None*, all Paths other than those described by [R69] **MUST** be ordered based on the Performance Metric specified in *primary*, from best to worst.

[D5] When the value of the PERFORMANCE Policy Criterion is not *None*, forwarding of IP Packets in an Application Flow to which the PERFORMANCE Policy Criterion is applied **SHOULD** favor Paths that appear earlier in the list of Paths that result from evaluation of the PERFORMANCE Policy Criterion.

A Path is selected for an Application Flow based on the *threshold* values as described in [R69], [R70], and [D5]. This Path selection can be reevaluated periodically, based on changes in the measured performance of the current Path, or based on other factors determined by the Service Provider. This reevaluation may or may not result in a new Path selection for the Application Flow. The examples in Appendix C demonstrate the effect of these requirements.

Once a Path is selected for an Application Flow it may be desirable to avoid frequent Path changes. Some Application Flows don't react well to frequent Path switching and the packet reordering and delay variation that can result. Similarly, moving several Application Flows from one Path to another can result in ping-ponging between Paths. The *ceiling* element provides a method to control this.

- [R71] If the value of the *threshold* element for the *primary* Performance Metric is *None*, the values of the *ceiling* and *remediation* elements **MUST** also be *None*.
- [R72] If the value of the *threshold* element for the *primary* Performance Metric is not *None*, the value of the *ceiling* element for the *primary* Performance Metric **MUST** be greater than or equal to the value of the *threshold* element for the *primary* Performance Metric or *None*.
- [R73] If the value of the *ceiling* element is not *None* and the measured value of the *primary* Performance Metric on the current Path is greater than or equal the *ceiling* element, the Service Provider **MUST** select a better Path for the Application Flow if one is available.
- [O1] For an Application Flow, if the measured value of the *primary* Performance Metric of the current Path is less than the *ceiling* value or if the *ceiling* value is *None*, the Service Provider **MAY** select a different Path that meets the constraints specified by [R69], [R70], and [D5].

To avoid excessive Path switching, the *ceiling* element allows for a range of measured performance values (between *threshold* and *ceiling*) within which the Service Provider is not obliged to attempt to switch the Application Flow to a different Path (see Figure 16). However, requirement [O1] indicates that the Service Provider can move an Application Flow to another Path at any time regardless of the measured performance of the current Path.

Another way to avoid moving an Application Flow to another Path is to attempt to remediate the performance on the current Path for the Application Flow. Remediation is a set of techniques intended to improve the performance for an Application Flow usually at the cost of increased resource usage. Techniques used for remediation are outside the scope of this document.

The *remediation* element for the *primary* Performance Metric indicates whether the Service Provider should perform remediation for this Application Flow and, if so, what performance level should trigger remediation.

- [R74] Application Flow performance remediation for the *primary* Performance Metric **MUST NOT** be performed if the value of the *remediation* element is *None* or if the measured performance of the current Path without remediation for the *primary* Performance Metric is less than the *remediation* value.

Note that [R74] makes the need for remediation conditional on the measured performance of the Path without remediation. This is important in order to avoid starting and stopping remediation around the value of the *remediation* element.

- [R75] If the value of the *ceiling* element and the *remediation* element for the *primary* Performance Metric are both not *None*, the value of the *remediation* element for the primary Performance Metric **MUST** be less than the value of the *ceiling* element for the *primary* Performance Metric.
- [R76] If the value of the *ceiling* element for the *primary* Performance Metric is *None* and the value of the *remediation* element for the *primary* Performance Metric is not *None*, the value of the *remediation* element for the primary Performance Metric **MUST** be less than or equal to the value of the *threshold* element for the *primary* Performance Metric.
- [D6] If the value of the *remediation* element is not *None*, the Service Provider **SHOULD** perform remediation for the *primary* Performance Metric when the measured performance of the current Path without remediation is greater than or equal to the value of the *remediation* element.

Figure 16 illustrates the relationship between these three parameters. In this diagram *remediation* is shown to have a value less than *threshold*, but consistent with [R75] it can have any value less than *ceiling*.




Figure 16 – PERFORMANCE Policy Criterion Parameters

As the measured performance of the current Path degrades, it will cross the *remediation* value, if specified, and remediation can be initiated (see [D6]). In the region below the *ceiling* value the Service Provider may choose a different Path for the Application Flow, if appropriate (see [O1]), but this consideration becomes more acute when the measured performance exceeds the *threshold* value (see [D5]). When the measured performance (taking into account any remediation in effect) exceeds the *ceiling* value, the Service Provider is required to choose a new Path for the Application Flow that has better performance, if one is available (see [R73]).

- [O2] If the Ingress Policy assigned to an Application Flow does not include the PERFORMANCE Policy Criterion or the value of the PERFORMANCE Policy Criterion is *None*, then the Service Provider **MAY** choose whether and how to consider performance in the selection of a Path for the Application Flow.

Examples of the use of the PERFORMANCE Policy Criterion can be found in Appendix C.

9.12.2.10 BANDWIDTH Ingress Policy Criterion

The BANDWIDTH Policy Criterion provides a method to express the intended bandwidth requirements for an Application Flow, and the probability of packet delay or discard in the face of varying bandwidth contention for Underlay Connectivity Service resources.

The effect of applying the BANDWIDTH Policy Criterion to an Application Flow is to declare IP Packets in the Application Flow either conformant or non-conformant. This can also include delaying IP Packets in order to improve conformance, i.e., traffic shaping. Non-conformant IP Packets in the Application Flow are discarded in order to meet the Policy.

The value of the BANDWIDTH Policy Criterion is either:

- an unnamed Rate Limiter 2-tuple $\langle commit, limit \rangle$ (see section 7.15), or
- the *name* of a named Rate Limiter specified in the SWVC List of Rate Limiters Service Attribute (see section 9.9)

A Rate Limiter constrains the bandwidth of a Bandwidth Flow which is the set of Application Flows that are assigned to the Rate Limiter. For unnamed Rate Limiters, the Bandwidth Flow comprises exactly one Application Flow. For named Rate Limiters, the Bandwidth Flow comprises one or more Application Flows that share the bandwidth constraints of the Rate Limiter. See section 7.15 for an explanation of unnamed and named Rate Limiters.

[R77] If the Ingress Policy assigned to an Application Flow does not include the BANDWIDTH Policy Criterion, the effect of the Policy **MUST** be as if the Policy includes the BANDWIDTH Policy Criterion with value $\langle 0, None \rangle$.

[R77] indicates that if a Policy does not contain a BANDWIDTH Policy Criterion, Application Flows to which that Policy is assigned do not have any committed bandwidth nor a bandwidth limit. This, in effect, means that no Rate Limiter is imposed on the Bandwidth Flow.

For the purpose of determining conformance, the information rate of the Application Flow is computed over time intervals of length *irduration*, which is an element in the value of the SWVC Performance Time Intervals Service Attribute (section 9.7).

[R78] The Subscriber and Service Provider **MUST** agree on sufficient Underlay Connectivity Service capacity to ensure that the *commit* values for all Bandwidth Flows can be met.

The agreement required by [R78] should consider other traffic that shares the capacity of the UCS such as traffic from other UNIs and traffic routed through the SD-WAN Edge. How the average information rates are determined and the behavior of the rate limiting function are described below.

This document does not specify the implementation of the bandwidth measurement/limiting by the Service Provider.

[R79] Bandwidth metering **MUST** occur after it is determined that the IP Packet will be forwarded based on Policy.

The intent of [R79] is to ensure that bandwidth metering occurs as late as possible in the packet processing in order to avoid metering IP Packets that are subsequently discarded.

The following requirements define the behavior imposed by the BANDWIDTH Policy Criterion.

- [R80] The information rate for IP Packets in a Bandwidth Flow that are declared conformant over any time interval equal to *irduration* **MUST** be at least the lower of *commit* and the information rate for IP Packets in that Bandwidth Flow that pass the metering point over that time interval.
- [R81] IP Packets in a Bandwidth Flow **MUST** be declared non-conformant in order to ensure that the information rate for such packets over any time interval equal to *irduration* that are declared conformant is, at most, *limit*.
- [O3] If the information rate for IP Packets in a Bandwidth Flow that pass the metering point over any time interval equal to *irduration* is greater than *commit*, then IP Packets in the Bandwidth Flow **MAY** be declared non-conformant in order to ensure that the information rate over that time interval for IP Packets across all Bandwidth Flows at a UNI does not exceed the available capacity of the Underlay Connectivity Services at that UNI.
- [O4] If, given the traffic received for various Bandwidth Flows with various Policies applied, the traffic that needs to be forwarded over a given Underlay Connectivity Service over any time interval equal to *irduration* exceeds the available capacity of that Underlay Connectivity Service, IP Packets in the affected Bandwidth Flows **MAY** be declared non-conformant in Bandwidth Flows where the information rate for IP Packets in the Bandwidth Flow that pass the metering point over that time interval is greater than *commit*.

The requirements above ([R80], [R81], [O3], and [O4]) describe how packets are declared conformant to the BANDWIDTH Policy (i.e., delivered with a high probability of success) or non-conformant (subject to discard as noted in [R83]). Note that [O5] allows for shaping of the Bandwidth Flow (i.e., delaying some IP Packets in the Bandwidth Flow) prior to this determination, which may result in fewer IP Packets being declared non-conformant and therefore subject to discard than otherwise would be.

[R80] indicates that at least the *commit* rate must be declared conformant if that much traffic is presented, i.e., if the average rate is less than or equal to *commit*, all of the traffic is declared conformant. On the other hand, [R81] indicates that the Service Provider discards packets to ensure that the average rate is less than *limit*.

[O3] and [O4] deal with the case where the average rate is between *commit* and *limit*. They indicate that the Service Provider may drop packets if the average rate is in this range if that is necessary to not exceed the total amount of UCS bandwidth or to avoid congestion on any particular UCS.

Note that the above requirements specify constraints over any time interval of duration *irduration*—i.e., they require a ‘sliding window’. Constraining bandwidth using a fixed, recurring,

window can have the effect of declaring bursts of traffic “green” that are twice as large as intended, as described in MEF 23.2 [31] Appendix H.2.

When the total amount of traffic received at an SD-WAN UNI exceeds the available capacity of the associated Underlay Connectivity Services, some packets may need to be discarded, even though each individual Bandwidth Flow is operating below the maximum specified in the BANDWIDTH Policy Criterion. The business importance (see section 9.12.2.8) of each Application Flow in the Bandwidth Flows is considered when packets need to be discarded, so long as each Bandwidth Flow still gets its *commit* rate.

Similarly, even when the total amount of traffic is less than the total available capacity on the Underlay Connectivity Services, as a result of other Policy Criteria affecting Path selection the combination of traffic across different Application Flows with different Policies might mean that the traffic that needs to be forwarded over a given Underlay Connectivity Service exceeds its capacity. As above, the business importance of the Application Flows is considered when packets need to be discarded.

[R82] IP Packets in a Bandwidth Flow **MUST** be declared conformant unless it is declared non-conformant by a condition specified in [R81], [O3], or [O4].

[O5] The Service Provider **MAY** delay certain IP Packets in a given Bandwidth Flow before determining conformance in order to increase the number of IP Packets in the Bandwidth Flow that are declared conformant.

The intent of [O5] is to note that the Service Provider may include a shaping function as part of the overall rate limiting process in order to potentially increase the number of IP Packets that are declared conformant. If shaping is performed, it can have an impact on the delay and delay variation of the affected Application Flows.

[R83] IP Packets in a Bandwidth Flow that are declared non-conformant **MUST** be discarded.

9.12.2.11 AF-SECURITY-INGRESS Ingress Policy Criterion

The AF-SECURITY-INGRESS Ingress Policy Criterion is used to select Application Flow Security Functions at the Ingress UNI as described in MEF 88 [38]. The value of this Policy Criterion is *None* or an Application Flow Security Policy Identifier from the value of the SWVC List of Security Policies Service Attribute (see section 9.9).

[R84] If the Ingress Policy assigned to an Application Flow does not include the AF-SECURITY-INGRESS Policy Criterion, the effect of the Policy **MUST** be as if the Policy includes the AF-SECURITY-INGRESS Policy Criterion with value *None*.

9.12.3 Egress Policy Criteria

This section describes the parameters and behavior of the EGRESS-BLOCK and AF-SECURITY-EGRESS Egress Policy Criteria.

9.12.3.1 EGRESS-BLOCK Egress Policy Criterion

The EGRESS-BLOCK Policy Criterion indicates whether the Egress Application Flow can egress at the UNI. It can have values *Allow* and *Discard*.

- [R85] If the value of the EGRESS-BLOCK Policy Criterion is *Discard*, all IP Packets in the Application Flow **MUST** be discarded.
- [R86] If the Egress Policy assigned to an Application Flow does not include the EGRESS-BLOCK Policy Criterion, then the behavior **MUST** be as if the value of the EGRESS-BLOCK Policy Criterion were *Allow*.

If the value of this Policy Criterion is *Allow*, then this Policy Criterion has no effect on the Egress Application Flow. Other Egress Policy Criteria may have an effect on the forwarding of the Application Flow.

9.12.3.2 AF-SECURITY-EGRESS Egress Policy Criterion

The AF-SECURITY-EGRESS Egress Policy Criterion is used to select Application Flow Security Functions at the SWVC End Point associated with an Egress UNI, as described in MEF 88 [38]. The value of this Policy Criterion is *None* or an Application Flow Security Policy Identifier from the value of the SWVC List of Security Policies Service Attribute (see section 9.9).

- [R87] If the Egress Policy assigned to an Application Flow does not include the AF-SECURITY-EGRESS Policy Criterion, the effect of the Policy **MUST** be as if the Policy includes the AF-SECURITY-EGRESS Policy Criterion with value *None*.

9.13 SWVC List of Application Flow Specification Groups Service Attribute

The value of the SWVC List of Application Flow Specification Groups Service Attribute is a list (possibly empty) of Application Flow Specification Group names.

An Application Flow Specification (see sections 8 and 9.14) can be a member of an Application Flow Specification Group. Application Flow Specification Groups provide a mechanism for associating a Policy with multiple Application Flows. If a Policy is associated with an Application Flow Specification Group and Zone, then that Policy is associated with all Application Flows that match Application Flow Specifications (in that Zone) that are members of the Application Flow Specification Group. An Application Flow Specification that is a member of the Application Flow Specification Group could, nonetheless, have an explicit Policy assigned at the SWVC End Point (in the same Zone) in which case that assignment would supersede the Application Flow Specification Group assignment for that Application Flow Specification.

For example, Application Flow Specifications *banana*, *pear*, and *grape* can all be members of the Application Flow Specification Group *fruit*. Group membership is indicated in the definition of each Application Flow Specification (see section 9.14).

- [R88] Each Application Flow Specification Group name in the value of the SWVC List of Application Flow Specification Groups Service Attribute **MUST** be an Identifier String.
- [R89] Each Application Flow Specification Group name in the value of the SWVC List of Application Flow Specification Groups Service Attribute **MUST** appear, at most, once.
- [R90] An Application Flow Specification Group name in the value of the SWVC List of Application Flow Specification Groups Service Attribute **MUST NOT** have the value *None*.

9.14 SWVC List of Application Flow Specifications Service Attribute

The SWVC List of Application Flow Specifications Service Attribute specifies the Application Flows that can be recognized by the SD-WAN service and information about how to identify IP Packets in each Application Flow. The value of the Service Attribute is a non-empty ordered list of 3-tuples $\langle AFName, AFCritList, AFGroup \rangle$ where:

- *AFName* is an Identifier String that is used to refer to the Application Flow Specification (and resulting Application Flows).
- *AFCritList* is a non-empty list of Application Flow Criteria 2-tuples of the form $\langle AFCritName, AFCritValue \rangle$ where:
 - *AFCritName* is an Identifier String containing an Application Flow Criterion Name from Table 8 or Table 9.
 - *AFCritValue* contains the parameter values specific to the Application Flow Criterion specified in *AFCritName*. If there are no parameter values, *AFCritValue* is *None*.
- *AFGroup* is an Application Flow Specification Group name contained in the value of the SWVC List of Application Flow Specification Groups Service Attribute (section 9.13) or *None* if the Application Flow is not a member of an Application Flow Specification Group.

- [R91] Each Application Flow Specification name, *AFName*, in the value of the SWVC List of Application Flow Specifications Service Attribute **MUST** appear, at most, once.
- [R92] Each Application Flow Specification name, *AFName*, in the value of the SWVC List of Application Flow Specifications Service Attribute **MUST NOT** be the same as an Application Flow Specification Group Name in the value of the SWVC List of Application Flow Specification Groups Service Attribute.
- [R93] An *AFCritName* **MUST NOT** appear more than once in the *AFCritList* for each item in the value of the SWVC List of Application Flow Specifications Service Attribute.
- [R94] If the *AFCritList* element in an entry of the SWVC List of Application Flow Specifications Service Attribute contains more than one Application Flow

Criterion, an Ingress IP Packet **MUST** match all Application Flow Criteria in order to be associated with the Application Flow.

- [R95] Each Ingress IP Packet **MUST** be assigned to an Application Flow based on the first Application Flow Specification in the value of the SWVC List of Application Flow Specifications Service Attribute whose Application Flow Criteria it matches, if any.
- [R96] Any Ingress IP Packet that cannot be associated with an Application Flow based on the value of the List of Application Flow Specifications Service Attribute **MUST** be discarded.

As shown in the example later in this section, the criteria for one Application Flow Specification can be a subset of the criteria for another Application Flow Specification, so the order that the Application Flow Specifications are matched, and hence the order of the Application Flow Specification definitions in the value of this Service Attribute is important and is one aspect of the agreed value of this Service Attribute.

[R94] indicates that the Application Flow Specification is defined by the conjunction of a set of Application Flow Criteria. This doesn't allow for alternatives within an Application Flow. This constraint is partially mitigated by the fact that most of the Application Flow Criteria are ranges or lists of values. Also, an Application Flow Specification Group can provide alternatives. For example, one Application Flow Specification can have criteria X and Y, and a second Application Flow Specification can have criteria X and W. If the two Application Flow Specifications are put into an Application Flow Specification Group, a common Policy can be applied to the Application Flow Specification Group, so it appears (almost) like a single Application Flow Specification defined as (X AND Y) OR (X AND W).

- [R97] The Application Flow Criteria supported by the Service Provider **MUST** include the Application Flow Criteria listed in Table 8.

<i>AFCritName</i>	<i>Layer</i>	<i>Match</i>	<i>Values for AFCritValue</i>	<i>Reference</i>
SAV4	3	IPv4 Source Address	List of IPv4 prefixes	RFC 791 [6]
DAV4	3	IPv4 Destination Address	List of IPv4 prefixes	RFC 791 [6]
PROTV4	3	IPv4 Protocol List	List of integers in the range 0 to 255 or a list of keywords from [1] or a mix of integers and keywords	IANA Protocol Numbers Registry [1]
SAV6	3	IPv6 Source Address	List of IPv6 prefixes	RFC 8200 [28]
DAV6	3	IPv6 Destination Address	List of IPv6 prefixes	RFC 8200 [28]
NEXTHEADV6	3	IPv6 Next Header List	List of integers in the range 0 to 255 or a list of keywords from [1] or a mix of integers and keywords	IANA Protocol Numbers Registry [1]
DSCP	3	Differentiated Services Code Point	List of integers in the range 0 to 63	RFC 2474 [12]

AFCritName	Layer	Match	Values for AFCritValue	Reference
SPORT	4	Transport Source Port ¹⁰	List of integers in the range 0 to 65535 or a list of service names from [2] or a mix of integers and service names	IANA Service Name and Port Number Registry [2]
DPORT	4	Transport Destination Port	List of integers in the range 0 to 65535 or a list of service names from [2] or a mix of integers and service names	IANA Service Name and Port Number Registry [2]
APPID	3-7	Custom match including heuristic/algorithmic matching	List of arguments starting with the Application Identifier.	This provides the ability to reference a library of custom application flow specifications.
ANY	1-7	Match any IP Packet	None	

Table 8 – Application Flow Criteria – Support Required

The Application Flow Criteria listed in Table 8 represent the basic “IP 5-tuple” (and APPID and ANY). Note that the IPv4 criteria are optional if all the UNIs that have an SWVC End Point for this SWVC only support IPv6 addressing, and the IPv6 criteria are optional only if those UNIs only support IPv4 addressing (see sections 11.4 and 11.5).

The APPID Application Flow Criterion provides a method for referring to named packet matching definitions (both simple and complex) defined by the Service Provider. These can include standard matches available to all the Service Provider’s Subscribers from a catalog and/or custom matches developed by the Service Provider by agreement with a particular Subscriber.

APPID can include simple protocol matches that can be accomplished with the other Policy Criteria, such as DPORT, (e.g., APPID match-dest-port 443) or *SSH* or *SNMP* or *RTP*, but they can also support more complex packet inspection, enabling identification of specific applications such as *Microsoft365*^{®11}, *Webex*^{®12}, or *Facebook*^{®13}.

- [R98]** If the Service Provider defines a named packet matching definition (either standard or custom) for use with the APPID Application Flow Criterion, the description provided to the Subscriber **MUST** include the following information:
- The Application Identifier
 - Additional Arguments Required (beyond the Identifier)
 - Details of the applications and application traffic matched by the APPID

¹⁰ This (and the other Application Flow Criteria referring to ports) was changed from “TCP/UDP Source Port” to “Transport Source Port” in MEF 70.1. This change allows the field to cover additional transport protocols (specified in PROTV4 or NEXTHEADV6) such as SCTP [23] and DCCP [21].

¹¹ Microsoft365 is a registered trademark of Microsoft Corporation

¹² Webex is a registered trademark of Cisco Systems Incorporated

¹³ Facebook is a registered trademark of Meta Platforms, Inc.

Complex matches, for example, using deep packet inspection, often require inspection of several initial packets and may include heuristics to define the characteristics of an Application Flow. These details are included in the description of the matching logic required by [R98].

For example:

- An APPID with name SIP: There are no additional arguments required, and the match is performed by inspecting the TCP or UDP source and destination port in each IP Packet for value 5060 or value 5061.
- An APPID named SIPUSER: This includes an additional argument “user-id”. The operation of this match is the same as SIP with the addition that if the port match is successful, the SIP *To* and *From* fields are matched against the “user-id”.

The Application Flow Criterion ANY matches all IP Packets. This criterion allows an Application Flow Specification to be defined that includes all “unmatched” IP Packets. A Policy can then be assigned to resulting Application Flows (per Zone) at the SWVC End Point. In general, if this Application Flow Criterion is used, it should be in the last Application Flow Specification in the list since no IP Packets are matched against subsequent Application Flow Specifications. An example is provided later in this section.

[R99] If an entry in the value of the SWVC List of Application Flow Specifications Service Attribute includes the Application Flow Criterion ANY, that entry **MUST NOT** contain any other Application Flow Criteria.

Support for the Application Flow Criteria listed in Table 9 is recommended. They provide additional capabilities or additional expressivity in Application Flow Specifications.

[D7] The Application Flow Criteria supported by the Service Provider **SHOULD** include the Application Flow Criteria listed in Table 9.

AFCritName	Layer	Match	Values for AFCritValue	Reference	Note
ETHERTYPE	2	EtherType	List of Integers in the range 0x0600 to 0xffff, e.g., 0x0800 for IPv4	802.3 [4]	Since SD-WAN is a layer 3 service, not all implementations have access to the L2 header during Application Flow matching, hence this Application Flow Criterion is optional.
SDAV4	3	IPv4 Source or Destination Address	List of IPv4 prefixes	RFC 791 [6]	This Application Flow Criterion allows an Application Flow to match a list of values for <i>either</i> the source or destination address. Support for this Application Flow Criterion is optional since it can be accomplished with two Application Flows based on the required criteria.

AFCritName	Layer	Match	Values for AFCritValue	Reference	Note
SDAV6	3	IPv6 Source or Destination Address	List of IPv6 prefixes	RFC 8200 [28]	Same as previous
SDPORT	4	Transport Source or Destination Port List	List of integers in the range 0 to 65535 or a list of service names from [2] or a mix of integers and service names	IANA Service Name and Port Number Registry [2]	This Application Flow Criterion allows an Application Flow to match a list of values for <i>either</i> the source or destination port. Support for this Application Flow Criterion is optional since it can be accomplished with two Application Flows based on the required criteria.

Table 9 – Application Flow Criteria – Support Recommended

It is important to note that Table 8 and Table 9 are not intended to describe the implementation or “syntax” of how Application Flows are described in any product or service, but rather the capabilities that are available. For example, no implementation is required to have a command or configuration parameter called DPORT, but rather that the implementation can “match packets with destination port =x”.

Here is an example value for this Service Attribute with four Application Flow Specifications:

```
[ {peach,      [{SAV4, [192.168.7.0/24] }, {DPORT, [80,443,8080] }], round }
  {VOIP,      [{APPID, ["RTP"] }], None }
  {banana,    [{DPORT, [80]}], long }
  {Else,      [{ANY, None}], None }
]
```

In this example, Application Flow Specification *peach* matches IP Packets from any 192.168.7.x address destined to port 80 or 443 or 8080 (and this Application Flow Specification is in the Application Flow Specification Group *round*). Application Flow Specification *VOIP* selects IP Packets that are matched by the APPID “RTP”. Application Flow Specification *banana* matches any IP Packet to port 80 that is not matched by *peach*, and this Application Flow Specification is in Application Flow Specification Group *long*. At the end of the list is the Application Flow Specification *Else*, which matches IP Packets not matched by the other three.

In this example, it is important that *banana* is after *peach* because there are some IP Packets that would match both definitions, but the desired behavior is that they are assigned to *peach*. If *banana* were first, then IP Packets with a source address in 192.168.7.x and destination port 80 would match to Application Flow Specification *banana* rather than Application Flow Specification *peach*.

It is also important to note that many Application Flow Specifications are not symmetric. For example, *banana* above, matches traffic destined *to* port 80, but response traffic will be *from* port 80. It can be made symmetric by using SDPORT, if that is available, or by adding another

Application Flow Specification, e.g., *rbanana*, which has SPORT 80. At each SD-WAN Edge a Policy can be assigned to one or both (with appropriate Zone qualification).

10 SD-WAN Virtual Connection (SWVC) End Point Service Attributes

The SWVC End Point is the construct that represents the attachment of an SWVC to a UNI. The SWVC End Point provides a container for attributes of the SWVC that can differ at each UNI.

This section describes Service Attributes at each SWVC End Point which are summarized in the following table, and each is described in more detail in the subsequent sections.

Attribute Name	Summary Description	Possible Values
SWVC End Point Identifier	Identification of the SWVC End Point for management purposes	Unique Identifier String for a given SWVC End Point.
SWVC End Point Associated UNI	Identifies the UNI that the SWVC End Point is associated with	An SD-WAN UNI Identifier
SWVC End Point List of UCS End Points	A list of the UCS End Points accessible by the SWVC End Point	List of UCS End Point Identifiers
SWVC End Point Policy Map	Maps Policies to Application Flows	A list of 4-tuples <zone, app, ipol, epol>
SWVC End Point Network Address Translation	Enables NAT/NAPT for Application Flows to and from the Internet	Disabled or 3-tuple <ipversion, publicprefix, exceptions>

Table 10 – Summary of SWVC End Point Service Attributes

10.1 SWVC End Point Identifier Service Attribute

The value of the SWVC End Point Identifier Service Attribute is a string that is used to allow the Subscriber and Service Provider to uniquely identify the association of the SWVC with a UNI.

[R100] The value of the SWVC End Point Identifier Service Attribute **MUST** be an Identifier String.

[R101] The value of the SWVC End Point Identifier Service Attribute **MUST** be unique across all SWVC End Points in the Service Provider Network.

10.2 SWVC End Point Associated UNI Service Attribute

The value of the SWVC End Point Associated UNI Service Attribute is an SD-WAN UNI Identifier Service Attribute value per section 11.1, which serves to specify the UNI that the SWVC End Point is associating with the SWVC. The SWVC End Point is said to be at this UNI.

10.3 SWVC End Point List of UCS End Points Service Attribute

Each SWVC End Point has a list of one or more UCS End Points to which it can forward Ingress IP Packets (from the UNI) and from which it can forward Egress IP Packets (towards the UNI). The value of the SWVC End Point List of UCS End Points Service Attribute is a non-empty list of UCS End Point Identifiers (section 12.3.1).

- [R102]** For any two SWVC End Points in an SWVC, one of the following **MUST** be true:
- the values of their respective SWVC End Point List of UCS End Points are identical
 - the intersection of their respective SWVC End Point List of UCS End Points is null

[R102] indicates that for every pair of SWVC End Points, either their List of UCS End Points are the same (have the same UCS End Points) which indicates that both are in the same SD-WAN Edge, or they have no UCS End Points in common which indicates that they are in different SD-WAN Edges. This means that every UNI in an SD-WAN Edge can use every UCS End Point in that SD-WAN Edge even if it only requires forwarding to a subset of them. Any desired restrictions on forwarding can be achieved by Policies.

10.4 SWVC End Point Policy Map Service Attribute

The SWVC End Point Policy Map specifies the Policies that are assigned to Application Flows at the SWVC End Point. The value of the SWVC End Point Policy Map is a (possibly empty)¹⁴ list of 4-tuples $\langle zone, app, ipol, epol \rangle$ where:

- *zone* is a Zone name.
- *app* is an Application Flow Specification or Application Flow Specification Group name.
- *ipol* is *None* or the name of an Ingress Policy.
- *epol* is *None* or the name of an Egress Policy.

The 4-tuple argument can be thought of as having two components. The pair $\langle zone, app \rangle$ defines the Application Flow (or Application Flows if *app* is an Application Flow Specification Group) at the UNI where this SWVC End Point is located, and the pair $\langle ipol, epol \rangle$ specify the Policy assignment to the Application Flow(s).

- [R103]** The *zone* element in each entry in the value of the SWVC End Point Policy Map **MUST** be one of the Zones listed in the SWVC List of Zones Service Attribute (section 9.5) or *Internet*.

- [R104]** The *app* element in each entry in the value of the SWVC End Point Policy Map **MUST** be either:
- an Application Flow Specification Group name from the value of the SWVC List of Application Flow Specification Groups Service Attribute (see section 9.13), or,
 - an Application Flow Specification name from the value of the SWVC List of Application Flow Specifications Service Attribute (see section 9.14)

¹⁴ An empty list only makes sense if all allowed Application Flows are covered by Zone-wide Policies and no Egress Policies are used.

- [R105] For a specific value of the *zone* element and a specified value of *app* in the value of the SWVC End Point Policy Map, the combination $\langle zone, app \rangle$ **MUST** appear in, at most, one entry in the list.
- [R106] The *ipol* element in each entry in the value of the SWVC End Point Policy Map **MUST** be either *None*, *Block*, or the name of an Ingress Policy from the value of the SWVC List of Policies Service Attribute (see section 9.12).
- [R107] If the *zone* element of an entry in the value of the SWVC End Point Policy Map is *Internet*, then the *ipol* element in that entry **MUST** be *None*.
- [R108] The *epol* element in each entry in the value of the SWVC End Point Policy Map **MUST** be either *None* or the name of an Egress Policy from the value of the SWVC List of Policies Service Attribute (see section 9.12).

10.4.1 Ingress Policy Assignment

An Ingress Policy can be assigned to an Application Flow either via the SWVC End Point Policy Map Service Attribute or via a Zone-wide Policy assignment specified in the SWVC List of Zones Service Attribute.

The SWVC List of Zones Service Attribute can assign a Zone-wide Ingress Policy to one or more Application Flows by associating a Policy with an Application Flow Specification (meaning the Policy is assigned to a single Application Flow at each UNI) or an Application Flow Specification Group (meaning the Policy is assigned to one or more Application Flows at each UNI). The Zone is implicit in both cases.

The first two elements $\langle zone, app \rangle$ of the value of the SWVC End Point Policy Map Service Attribute specify one or more Application Flows. The *app* element can be an Application Flow Specification, (thus specifying a single Application Flow at this UNI), or it can be an Application Flow Specification Group (thus specifying one or more Application Flows at this UNI). The *ipol* element identifies the Ingress Policy that is assigned to the specified Application Flows.

This results in four ways that an Ingress Policy could be assigned to an Application Flow identified by a given UNI, Application Flow Specification, and Zone:

- By using the SWVC List of Zones Service Attribute to associate a Policy with the Zone and the Application Flow Specification.
- By using the SWVC List of Zones Service Attribute to associate a Policy with the Zone and with an Application Flow Specification Group of which the Application Flow Specification is a member.
- By using the SWVC End Point Policy Map Service Attribute for the SWVC End Point at the UNI to associate a Policy with the Zone and the Application Flow Specification.
- By using the SWVC End Point Policy Map Service Attribute for the SWVC End Point at the UNI with the Zone and with an Application Flow Specification Group of which the Application Flow Specification is a member.

There is a well-defined precedence among the four. [R109], [R110], and [R111] define this precedence.

- [R109] Assignment of an Ingress Policy to an Application Flow via the SWVC End Point Policy Map Service Attribute **MUST** take precedence over assignment via the SWVC List of Zones Service Attribute.
- [R110] Assignment of an Ingress Policy to an Application Flow via an entry in the SWVC End Point Policy Map Service Attribute that specifies the corresponding Application Flow Specification **MUST** take precedence over assignment via an entry that specifies an Application Flow Specification Group of which the corresponding Application Flow Specification is a member.
- [R111] Assignment of an Ingress Policy to an Application Flow via an entry in the SWVC List of Zones Service Attribute that specifies the corresponding Application Flow Specification **MUST** take precedence over assignment via an entry that specifies an Application Flow Specification Group of which the corresponding Application Flow Specification is a member.

Note that [R109] means that a Policy assigned via the SWVC List of Zones Service Attribute might be overridden by a Policy assigned via the SWVC End Point Policy Map Service Attribute at some UNIs but not others. Requirements [R109] through [R111] describe a precedence for Policy assignment that is depicted in Figure 17.




Figure 17 – Precedence for Policy Assignment

- [R112] If an Application Flow is not assigned an Ingress Policy via any of the methods described in [R109] through [R111] at a UNI, Ingress IP Packets mapped to the Application Flow **MUST** be discarded.
- [R113] If an Application Flow is assigned the reserved Ingress Policy *Block* at a UNI, Ingress IP Packets mapped to the Application Flow **MUST** be discarded.

As indicated in the SD-WAN List of Policies Service Attribute (see section 9.12) the Policy name *Block* is reserved for the purpose indicated in [R113] and cannot be explicitly defined.

If an Ingress Policy is associated with an Application Flow Specification Group (in a Zone), then that Ingress Policy is assigned to all of the Application Flows that match that Zone and any of the Application Flow Specifications in the Application Flow Specification Group (assuming no higher precedence assignment). [R113] provides a mechanism to selectively disallow forwarding of one or more of those Application Flows by assigning the Ingress Policy *Block* (see section 9.12) to one

or more Application Flow Specifications in the Zone (since [R112] cannot be invoked). Note that [R38] includes additional conditions for forwarding an Application Flow.

Not forwarding an Application Flow (either due to not having an Ingress Policy assigned or by explicit assignment of the reserved Policy, *Block*), only refers to Application Flows that ingress at the UNI.

10.4.2 Egress Policy Assignment

An Egress Policy can be assigned to an Application Flow via the SWVC End Point Policy Map Service Attribute at the Egress SWVC End Point.

The first two elements $\langle zone, app \rangle$ of the value of the SWVC End Point Policy Map Service Attribute specify one or more Application Flows. The *app* element can be an Application Flow Specification, (thus specifying a single Application Flow at this UNI), or it can be an Application Flow Specification Group (thus specifying one or more Application Flows at this UNI). The *epol* element identifies the Egress Policy that is assigned to the specified Application Flows. Egress Policies can be assigned to IP Packets from the Internet by specifying the reserved Zone name, *Internet*.

An Application Flow at an Egress UNI is, in the general case, an aggregation of IP Packets from multiple Ingress Application Flows. Consider the following example:

- at SWVC End Point 1 there is an Ingress Application Flow that matches Application Flow Specification *email* in Zone *corp*. Some of this Application Flow is forwarded to SWVC End Point 2 and some to SWVC End Point 3.
- at SWVC End Point 2 there is an Ingress Application Flow that matches Application Flow *email* in Zone *corp*. Some of this Application Flow is forwarded to SWVC End Point 1 and some to SWVC End Point 3.
- At SWVC End Point 3 the Egress Application Flow that matches Application Flow Specification *email* in Zone *corp* is the aggregate of the portion forwarded from SWVC End Point 1 and the portion forwarded from SWVC End Point 2.

Egress Application Flows are delivered to an Egress UNI unless blocked by an Egress Policy (including a Security Function as described in MEF 88 [38] via the AF-SECURITY-EGRESS Policy Criterion (Section 9.12.3.2)).

- [R114]** If an Egress Policy is not assigned to an Application Flow at an Egress UNI, then Egress IP Packets that are in the Application Flow **MUST** be forwarded to the Egress UNI.

[R114] can occur if no Policies are mapped to the Application Flow at the SWVC End Point, or if the *epol* element is *None*.

If it is necessary to block an Application Flow at an Egress UNI, this can be achieved by specifying an Egress Policy (*epol* element) that includes an EGRESS-BLOCK Policy Criterion.

[R115] Assignment of an Egress Policy to an Application Flow via an entry in the SWVC End Point Policy Map Service Attribute that specifies the corresponding Application Flow Specification **MUST** take precedence over assignment via an entry that specifies an Application Flow Specification Group of which the corresponding Application Flow Specification is a member.

10.4.3 Examples of Ingress Application Flows and Policy Assignment

Figure 18 depicts the precedence for assigning an Ingress Policy to an Application Flow. The SD-WAN Service has three Subscriber sites (A, B, and C) with a single Zone, *hr*, defined. The SWVC has four Application Flow Specifications: *apple*, *banana*, *celery*, *basil* (via the SWVC List of Application Flow Specifications Service Attribute in section 9.14). *Apple* and *banana* are in Application Flow Specification Group *fruit*. Application Flow Specifications *celery* and *basil* are not in any Application Flow Specification Group. Zone *hr* has a single Zone-wide Policy assignment (*basil* ← *sprinkle*).

The SWVC also has three Ingress Policies: *bake*, *cook*, and *sprinkle* (via the SWVC List of Policies Service Attribute in section 9.9). All SWVCs have a reserved Policy named *Block*.

Note: In the diagram and subsequent explanatory text, *xxxx_{yy}* should be interpreted as the Application Flow that matches Application Flow Specification *xxxx* and is in Zone *yy*.

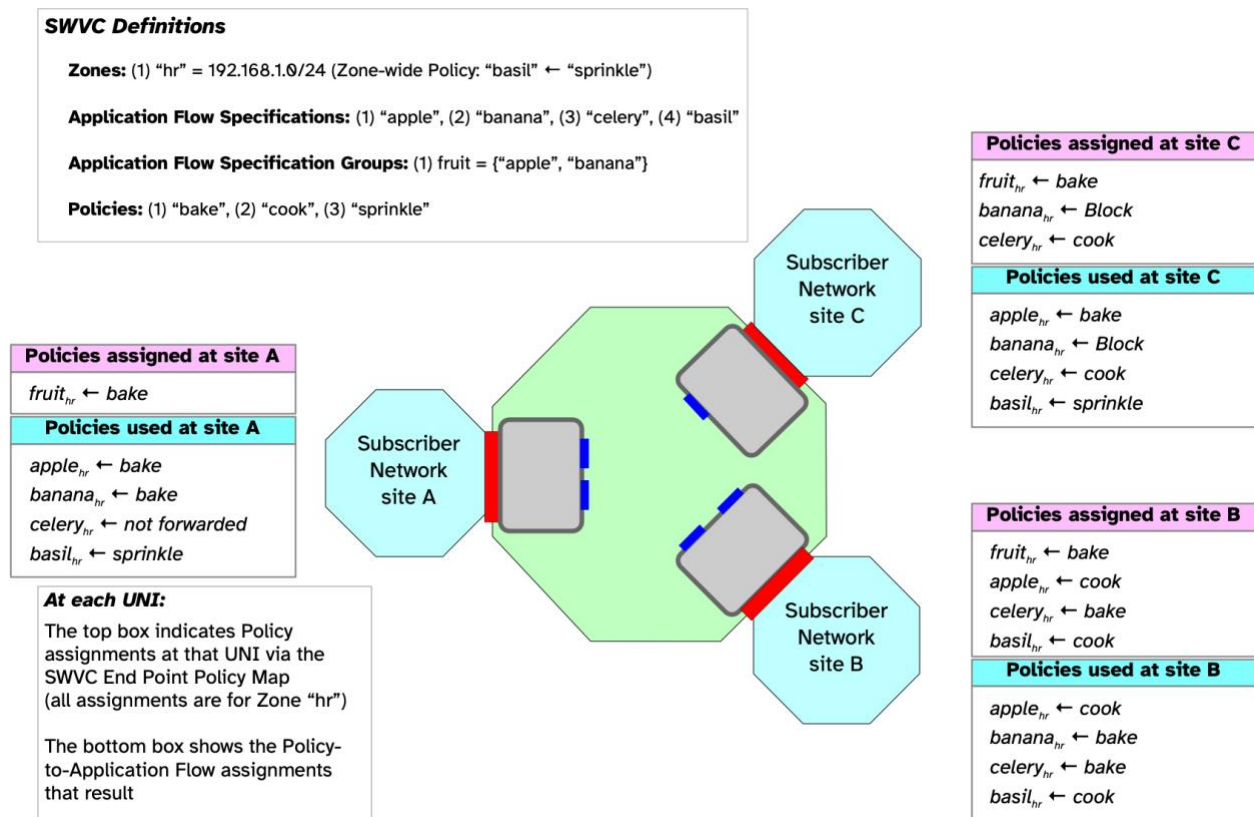


Figure 18 – Assigning Ingress Policies

At site A:

- Policy *bake* is associated with (via the SWVC End Point Policy Map) Application Flow Specification Group *fruit* in Zone *hr*. This means that *apple_{hr}* and *banana_{hr}* are forwarded based on Policy *bake*.
- Application flow *basil_{hr}* is forwarded based on the Zone-wide Policy *sprinkle*.
- Application flow *celery_{hr}* does not have a Zone-wide Policy nor is it assigned a Policy via the SWVC End Point Policy Map, so it is not forwarded.

At site B:

- Policy *bake* is associated with Application Flow Specification Group *fruit* in *hr*. So, *banana_{hr}* is forwarded based on Policy *bake*, but Policy *cook* is assigned to Application Flow *apple_{hr}*, overriding the Application Flow Specification Group association.
- Application Flow *celery_{hr}* is forwarded based on the Policy *bake* (via the SWVC End Point Policy Map).
- Application Flow *basil_{hr}* is forwarded based on the Policy *cook* (via the SWVC End Point Policy Map overriding the Zone-wide Policy).

At site C:

- Policy *bake* is associated with Application Flow Specification Group *fruit* in *hr*. So, *apple_{hr}* is forwarded based on Policy *bake*, but Policy *Block* is assigned to Application Flow *banana_{hr}*, so it is not forwarded.
- Application Flow *celery_{hr}* is forwarded based on the Policy *cook* (via the SWVC End Point Policy Map).
- Application Flow *basil_{hr}* is forwarded based on the Zone-wide Policy *sprinkle*.

10.5 SWVC End Point Network Address Translation Service Attribute

When Ingress IP Packets are destined to the Internet via Internet Breakout (see section 7.14), the Subscriber might require the SD-WAN Service Provider to perform Network Address Translation (NAT), translating the Subscriber's IP addresses into public IP addresses and vice versa.

This Service Attribute determines whether and how NAT is performed by the Service Provider in the SD-WAN Edge. NAT may also be performed by the UCS Service Provider within the UCS or by the Subscriber within the Subscriber Network. NAT performed in the UCS is determined by the UCS Service Attributes which are agreed between the UCS Service Provider and UCS subscriber.

The value of the SWVC End Point Network Address Translation Service Attribute is either *Disabled*, or a 3-tuple (*ipversion, publicprefix, exceptions*) where:

- *ipversion* is a string that indicates the type(s) of Subscriber addresses that are translated. *IPv4* indicates that IPv4 addresses are translated, *IPv6* indicates that IPv6 addresses are translated, and *Both* indicates that both IPv4 and IPv6 addresses are translated.
- *publicprefix* is a publicly assigned IPv4 Prefix representing the public address (if /32) or addresses that the Subscriber's IP addresses are translated to (and from).
- *exceptions* is a list (possibly empty) of IP prefixes that are not subject to translation.

Translation to IPv6 public addresses is outside the scope of this document.

- [R116]** If the value of the *ipversion* element of the SWVC End Point Network Address Translation Service Attribute is *IPv4* or *Both*, the following IP Packets **MUST** be subject to behavior consistent with Network Address Translation and Network Address Port Translation as described in RFC 3022 [14], using an IPv4 address in *publicprefix*:
- IPv4 Packets that ingress at the UNI associated with this SWVC End Point and transmitted towards the Internet via Internet Breakout with a source IP address not contained in any of the IP prefixes specified in *exceptions*
 - IPv4 Packets received from the Internet and destined to an address in *publicprefix*
- [R117]** If the value of the *ipversion* element of the SWVC End Point Network Address Translation Service Attribute is *IPv6* or *Both*, the following IP Packets **MUST** be subject to behavior consistent with Network Address Translation and Network Address Port Translation as described in RFC 3022 [14], using an IPv4 address in *publicprefix*:
- IPv6 Packets that ingress at the UNI associated with this SWVC End Point and transmitted towards the Internet via Internet Breakout with a source IP address not contained in any of the IP prefixes specified in *exceptions*
 - IPv4 Packets received from the Internet and destined to an address in *publicprefix*

Note that in the case of IPv6, the Subscriber IPv6 addresses are translated to IPv4 addresses, and thus in the reverse direction, IPv4 addresses are received from the Internet.

- [R118]** IP Packets that ingress at the UNI associated with this SWVC End Point and are transmitted towards the Internet via Internet Breakout that are not subject to NAT/NAPT as specified in [R116] and [R117] **MUST** be forwarded without address or port translation.

[R118] is intended to address the situation where either IPv4 or IPv6 packets are subject to translation, but not both types.

Note that the Service Provider might perform multiple stages of NAT; this is consistent with RFC 3022 [14] and hence compliant with [R116] and [R117] provided that the IP address eventually used to send/receive IP packets to/from the Internet is within the specified *publicprefix*. The IPv4 address used in any intermediate stages of NAT need not be from a publicly assigned IPv4 Prefix.

11 SD-WAN UNI Service Attributes

The UNI is a reference point that represents the demarcation between the responsibility of the Subscriber and the responsibility of the Service Provider. As a result, at any given UNI there is only a single Subscriber and a single Service Provider (see [R1] and [R2]).

This section includes the Service Attributes at each UNI which are summarized in the following table and described in more detail in the subsequent sections. Since an SD-WAN Service delivers IP packets between multiple Subscriber Locations, much of this section is adapted from the UNI Service Attributes and UNI Access Link Service Attributes section of the IP Service Attributes MEF Standard, MEF 61.1 [32] and MEF 61.1.1 [33] in order to achieve the greatest amount of commonality between MEF IP Services and MEF SD-WAN Services.

Attribute Name	Summary Description	Possible Values
SD-WAN UNI Identifier	Identification of the UNI for management purposes	Unique Identifier String for the UNI
SD-WAN UNI L2 Interface	Describes the underlying L2 interface for the UNI	<i>Untagged or CVLAN x</i>
SD-WAN UNI L2 Maximum Frame Size	Specifies the maximum length L2 frame that is accepted by the Service Provider	An integer number of bytes ≥ 1522
SD-WAN UNI IPv4 Connection Addressing	IPv4 Connection Address mechanism	<i>None, DHCP, or Static</i>
SD-WAN UNI IPv6 Connection Addressing	IPv6 Connection Address mechanism	<i>None, DHCP, SLAAC, Static or LL-only</i>
SD-WAN UNI Routing Protocols	List of Routing Protocols used across the UNI	See section 11.6
SD-WAN UNI Availability Objective	The goal for the lower bound on the availability of the UNI during each reporting period	<i>None</i> or a percentage between 0 and 100%
SD-WAN UNI Mean Time To Repair Objective	The goal for the upper bound on the average length of a UNI outage during each reporting period.	<i>None</i> or a number of minutes

Table 11 – Summary of SD-WAN UNI Service Attributes

11.1 SD-WAN UNI Identifier Service Attribute

The value of the SD-WAN UNI Identifier Service Attribute is a string that is used to allow the Subscriber and Service Provider to uniquely identify the UNI.

[R119] The value of the SD-WAN UNI Identifier Service Attribute **MUST** be an Identifier String.

[R120] The value of the SD-WAN UNI Identifier Service Attribute **MUST** be unique across all SD-WAN UNIs in the Service Provider Network.

As an example, the Subscriber and Service Provider might agree to use “NY-1” as a value of the SD-WAN UNI Identifier Service Attribute.

11.2 SD-WAN UNI L2 Interface Service Attribute

The SD-WAN UNI L2 Interface Service Attribute describes the underlying network layer that carries IP Packets across the UNI. The fundamental role of the UNI is to be able to convey IP Packets (layer 3) between the Subscriber and the SP.

The SD-WAN UNI layer 2 is an Ethernet MAC. The value of this Service Attribute describes the set of Ethernet MAC frames that are mapped to the UNI at layer 2. The possible values are *Untagged* and *CVLAN x* where *x* is a VLAN ID in the range 1 to 4094.

[R121] The format for an L2 frame that crosses the UNI **MUST** be that of the Ethernet MAC Frame that is specified in Clause 3 of IEEE Std 802.3-2018 [4] except for section 3.2.7 of that document.

Note that [R121] means that Ethernet MAC frames will be discarded by the SD-WAN Edge if they are not properly constructed. For example, a Service Frame with an incorrect Frame Check Sequence will be discarded. Section 3.2.7 in IEEE Std. 802.3-2018 [4] describes the maximum length of the client data field and limits it to 1982 bytes (2000-byte frame), however this specification does not impose this limit. The Subscriber and the Service Provider can agree to any value subject to the constraints described in the SD-WAN UNI L2 Maximum Frame Size Service Attribute (section 11.3).

The following Ethernet MAC Frame formats are defined:

- When the field following the Source Address field is a TPID (Tag Protocol ID defined in IEEE Std 802.1Q-2018 [3]) with the value 0x8100 and the corresponding VLAN ID is not 0x0000, the Ethernet MAC Frame is said to be a C-Tagged frame and the value of the VLAN ID field is referred to as the C-VID.
- When the two bytes following the Source Address do not contain the values 0x8100 or 0x88a8, the Ethernet MAC Frame is said to be an Untagged frame.
- When the field following the Source Address field is a TPID (Tag Protocol ID defined in IEEE Std 802.1Q-2018 [3]) with the value 0x8100 and the corresponding VLAN ID is 0x0000, the Ethernet MAC Frame is said to be a Priority Tagged frame and is treated as an Untagged frame.

C-Tagged and Untagged frames are eligible to be mapped to the UNI, subject to the constraints imposed by the value of this Service Attribute. Handling of other frame types including S-Tagged frames (the two bytes following the Source Address contain the value 0x88a8) is beyond the scope of this document.

[R122] If the value of the SD-WAN UNI L2 Interface Service Attribute is *Untagged*, then all Untagged Ethernet MAC Frames on the underlying physical or virtual Ethernet link **MUST** be delivered to the UNI.

[R123] If the value of the SD-WAN UNI L2 Interface Service Attribute is *Untagged*, then C-tagged Ethernet MAC Frames on the underlying physical or virtual Ethernet link **MUST NOT** be delivered to the UNI.

- [R124] If the value of the SD-WAN UNI L2 Interface Service Attribute is *CVLAN x*, then all C-tagged Ethernet MAC Frames with a C-VID of *x* on the underlying physical or virtual Ethernet link **MUST** be delivered to the UNI.
- [R125] If the value of the SD-WAN UNI L2 Interface Service Attribute is *CVLAN x*, then Untagged and C-tagged Ethernet MAC Frames with a C-VID other than *x* on the underlying physical or virtual Ethernet link **MUST NOT** be delivered to the UNI.

The details and parameters of the layer 1 interface, including the speed of the interface, must be agreed to between the Subscriber and the Service Provider but are beyond the scope of this specification.

The implication of supporting a single VLAN at a UNI (i.e., *CVLAN x*) is that the layer 1 channel that is conveying the Ethernet MAC frames can potentially access multiple UNIs. Each C-VID value can be mapped to a different UNI for the same SD-WAN Service, or a different SD-WAN Service, or, in theory, a different type of service. Details of this capability are out of scope for this document but may be included in a future version.

11.3 SD-WAN UNI L2 Maximum Frame Size Service Attribute

The SD-WAN UNI L2 Maximum Frame Size Service Attribute specifies the maximum Ethernet MAC frame size that will be accepted by the Service Provider at the UNI.

- [R126] The value for the SD-WAN UNI L2 Maximum Frame Size **MUST** be an integer number of bytes ≥ 1522 .
- [D8] Any L2 Frame that crosses the Ingress UNI whose length exceeds the value of the SD-WAN UNI L2 Maximum Frame Size **SHOULD** be discarded.

11.4 SD-WAN UNI IPv4 Connection Addressing Service Attribute

The SD-WAN UNI IPv4 Connection Addressing Service Attribute specifies how IPv4 addresses are allocated to the devices on the Subscriber side of the UNI. The Service Attribute has one of three possible values: *None*, *DHCP*, or *Static*. In the case of *DHCP* and *Static* there are some additional parameters.

If the IPv4 Connection Addressing is *None*, no IPv4 addresses are used and IPv4 is disabled on the link. Note that in this case IPv6 connection addresses are needed.

- [R127] If the value of the SD-WAN UNI IPv4 Connection Addressing Service Attribute is *None*, IPv4 Packets **MUST NOT** be forwarded to or from the UNI.
- [R128] The SD-WAN UNI IPv4 Connection Addressing Service Attribute and the SD-WAN UNI IPv6 Connection Addressing Service Attribute (section 11.5) **MUST NOT** both have the value *None*.

If the value of the SD-WAN UNI IPv4 Connection Addressing Service Attribute is *DHCP*, then DHCP is used by devices in the Subscriber Network to request IPv4 addresses in a given subnet from the Service Provider as described in RFC 2131 [8] and RFC 2132 [9]. The Service Provider hosts the DHCP server, and the Subscriber devices act as the DHCP clients.

- [R129]** When the IPv4 Connection Addressing is *DHCP*, the Service Provider **MUST** use DHCP to convey to the Subscriber, in addition to the IPv4 address, the subnet mask and the default router address.

If the value of the SD-WAN UNI IPv4 Connection Addressing is *Static*, then IPv4 addresses in a given IPv4 subnet are statically assigned to the Service Provider and the Subscriber.

For *DHCP* and *Static*, a number of further parameters have to be agreed including:

- Primary Subnet:
 - IPv4 Prefix (IPv4 address prefix and mask length between 0 and 31, in bits)
 - Service Provider IPv4 Addresses (Non-empty list of IPv4 addresses)
- Optional Secondary Subnet List; each entry containing:
 - IPv4 Prefix (IPv4 address prefix and mask length between 0 and 31, in bits)
 - Service Provider IPv4 Addresses (Non-empty list of IPv4 addresses)

The parameters consist of a primary subnet and zero or more secondary subnets. In each case, the IP Prefix is specified, along with the Service Provider's IPv4 addresses. In the case of the primary subnet, this IP Prefix is referred to as the Connection Primary IPv4 Prefix, and for a secondary subnet, the Connection Secondary IPv4 Prefix.

For *DHCP* the address of the Subscriber's default router is provided in the DHCP response. For *Static* addressing, the Service Provider's addresses are assumed to be the default router addresses.

Note that the IPv4 Prefix and Service Provider addresses need to be agreed even when DHCP is used, so that the Subscriber can ensure they do not conflict with any other addressing used within the Subscriber Network.

If DHCP is used, the IPv4 address range from which Subscriber addresses are dynamically assigned is taken from the prefixes listed in the value of the SWVC Reserved Prefixes Service Attribute that are subnets of the Connection Primary IPv4 Prefix or a Connection Secondary IPv4 Prefix.

- [R130]** If the value of SD-WAN UNI IPv4 Connection Addressing is *DHCP*, addresses that are dynamically assigned by DHCP within the Connection Primary IPv4 Prefix or a Connection Secondary IPv4 Prefix **MUST** be taken from within an IP Prefix listed in the SWVC Reserved Prefixes Service Attribute (section 9.4) that is a subnet of the Connection Primary IPv4 Prefix or Connection Secondary IPv4 Prefix.

For example, if the SWVC List of Reserved Prefixes includes:

- 192.168.1.0/26
- 192.168.2.0/26

and the Connection Primary IPv4 Prefix is 192.168.1.0/24, DHCP can dynamically assign addresses 192.168.1.1 through 63 and the Subscriber can assign addresses 192.168.1.64 through 254 (note that [R132] prohibits the Subscriber from assigning the highest address in the prefix).

[R131] If the value of the SD-WAN UNI IPv4 Connection Addressing is *Static* or *DHCP*, for the Primary Subnet and for each Secondary Subnet, the Service Provider IPv4 Addresses **MUST** be within the specified IPv4 Prefix.

The Subscriber can statically assign any IPv4 address within the subnets identified by the Connection IPv4 Prefixes, other than the Service Provider address itself, the lowest and highest possible addresses, which are generally reserved, and any addresses within IP Prefixes listed in the value of the SWVC Reserved Prefixes Service Attribute (see section 9.4).

[R132] If the value of the SD-WAN UNI IPv4 Connection Addressing is *DHCP* or *Static*, the Subscriber **MUST NOT** statically assign any of the following for use by Subscriber devices connected to the UNI:

- Any IPv4 address that is neither within the Connection Primary IPv4 Prefix nor within the Connection Secondary IPv4 Prefix specified in an entry in the Secondary Subnet List.
- Any of the Primary Subnet Service Provider IPv4 Addresses.
- Any of the Service Provider IPv4 Addresses specified in an entry in the Secondary Subnet List.
- The lowest and highest IPv4 addresses in the Connection Primary IPv4 Prefix, if the prefix length is less than or equal to 30.
- The lowest and highest IPv4 addresses in the Connection Secondary IPv4 Prefix specified in an entry in the Secondary Subnet List, if the prefix length is less than or equal to 30.
- Any IPv4 address within IP Prefixes listed in the value of the SWVC Reserved Prefixes Service Attribute (see section 9.4).

11.5 SD-WAN UNI IPv6 Connection Addressing Service Attribute

The SD-WAN UNI IPv6 Connection Addressing Service Attribute specifies how IPv6 addresses are allocated to the devices connected to the UNI. The Service Attribute has one five values *None*, *DHCP*, *SLAAC*, *Static* or *LL-only*, plus in the case of *DHCP*, *SLAAC* or *Static*, some additional parameters. If the IPv6 Connection Addressing is *None*, no IPv6 addresses are used by the devices connected to the UNI and IPv6 is disabled on the link. Note that in this case IPv4 connection addresses are needed (see [R128]).

[R133] If the value of the SD-WAN UNI IPv6 Connection Addressing Service Attribute is *None*, IPv6 Packets **MUST NOT** be forwarded to or from the UNI.

If the value of the SD-WAN UNI IPv6 Connection Addressing Service Attribute is one of *DHCP*, *Static*, *SLAAC*, or *LL-only*, then IPv6 Link-Local addresses are present on the UNI. If the value is *LL-only*, then only IPv6 Link-Local addressing is used on the UNI.

If the value of the SD-WAN UNI IPv6 Connection Addressing is *DHCP*, then DHCPv6 is used by the Subscriber devices to request IPv6 addresses in a given subnet from the Service Provider as described in RFC 3315 [18]. The Service Provider hosts the DHCP server, and the Subscriber devices act as the DHCP clients.

[R134] When the value of the SD-WAN UNI IPv6 Connection Addressing Service Attribute is *DHCP*, the Service Provider **MUST** use DHCP to convey to the Subscriber, in addition to the IPv6 address, the subnet mask and the default router address.

If the value of the SD-WAN UNI IPv6 Connection Addressing is *Static*, then IPv6 addresses in a given IPv6 subnet are statically assigned to the Service Provider and the Subscriber.

If the value of the SD-WAN UNI IPv6 Connection Addressing is *SLAAC*, then Stateless Address Autoconfiguration (SLAAC) is used by the Subscriber devices to create unique IPv6 global addresses within an IP Prefix advertised by the Service Provider as described in RFC 4862 [22]. The Router Advertisements that convey the IP Prefix are also used to convey the prefix length and router address.

For *DHCP*, *SLAAC* and *Static*, several further parameters have to be agreed:

- Subnet List of one or more subnets, each comprising:
 - IPv6 Prefix
 - IPv6 address prefix and prefix length between 0 and 128 for *DHCP* and *Static*, or
 - IPv6 address prefix and prefix length of 64 for *SLAAC*
 - Service Provider IPv6 Addresses (Non-empty list of IPv6 addresses)

The parameters consist of a list of one or more subnets. For each subnet, the IPv6 prefix and the Service Provider's IPv6 address are specified. The IPv6 Prefix is referred to as a Connection IPv6 Prefix. Note that an IP Prefix and Service Provider addresses need to be agreed even when DHCP or SLAAC is used, so that the Subscriber can ensure they do not conflict with any other addressing used within the Subscriber Network.

A list (possibly empty) of reserved IP Prefixes can be specified (section 9.4); these specify IP addresses that are not available for the Subscriber to assign statically.

If DHCP is used, the IPv6 address range, from which Subscriber addresses are dynamically assigned, is taken from the prefixes listed in the value of the SWVC Reserved Prefixes Service Attribute that are subnets of any Connection IPv6 Prefix.

[R135] If the value of the SD-WAN UNI IPv6 Connection Addressing is *Static*, *DHCP* or *SLAAC*, for each entry in the Subnet List, the Service Provider IPv6 Addresses **MUST** be within the Connection IPv6 Prefix for that entry.

- [R136] If the value of the SD-WAN UNI IPv6 Connection Addressing is *DHCP*, addresses that are dynamically assigned by DHCP within one of the Connection IPv6 Prefixes **MUST** be taken from within an IP Prefix in the value of the SWVC Reserved Prefixes (section 9.4) that is a subnet of that Connection IPv6 Prefix.
- [R137] If the value of the SD-WAN UNI IPv6 Connection Addressing is *SLAAC*, the IP Prefix advertised by the Service Provider as described in RFC 4862 [22] using Router Advertisements **MUST** be the Connection IPv6 Prefix for the first entry in the Subnet List.

The Subscriber can statically assign any IPv6 address within the subnets identified by the Connection IPv6 Prefix in each entry, other than the Service Provider address itself, the lowest and highest possible addresses, which are generally reserved, and any address within IP Prefixes listed in the value of the SWVC Reserved Prefixes Service Attribute (see section 9.4).

- [R138] If the value of the SD-WAN UNI IPv6 Connection Addressing is *DHCP*, *SLAAC* or *Static*, the Subscriber **MUST NOT** statically assign any of the following for use on the UNI by Subscriber devices:
- Any IPv6 address that is not within the Connection IPv6 Prefix for an entry in the Subnet List.
 - Any IPv6 address within the Connection IPv6 Prefix for the first entry in the Subnet List, if the SD-WAN UNI IPv6 Connection Addressing is *SLAAC*.
 - Any of the Service Provider IPv6 Addresses specified in an entry in the Subnet List.
 - The lowest and highest IPv6 addresses in the Connection IPv6 Prefix for an entry in the Subnet List if the prefix length is less than or equal to 126.
 - Any IPv6 address within IP Prefixes listed in the value of the SWVC Reserved Prefixes Service Attribute (see section 9.4).

11.6 SD-WAN UNI Routing Protocols Service Attribute

The SD-WAN UNI Routing Protocols Service Attribute specifies the routing protocols and associated parameters that are used to exchange IP routes across the UNI. The value is a list of protocols (possibly empty), where each entry consists of the protocol name (one of *Static*, *BGP*, or *OSPF*), the type of routes that will be exchanged (one of *IPv4*, *IPv6* or *Both*), and a set of additional parameters as specified in the subsections below.

- [R139] The value of the SD-WAN UNI Routing Protocols Service Attribute **MUST NOT** contain more than one entry for the same protocol name, except when there are exactly two entries with a given protocol name, one with route type IPv4 and one with route type IPv6.

Note that regardless of the routing protocol in use, the Service Provider is responsible for forwarding IP Packets with a destination IP Address within the IP Prefixes identified by the SD-WAN UNI IPv4 Connection Addressing Service Attribute (see section 11.4) and the SD-WAN

UNI IPv6 Connection Addressing Service Attribute (see section 11.5) toward the corresponding UNI.

When all of the end hosts in the Subscriber Network that are reachable at a given UNI are directly adjacent (at Layer 3) to that UNI (i.e. there is no router on the Subscriber's side of the UNI), and therefore only use IP addresses within the IP Prefixes identified by the SD-WAN UNI IPv4 Connection Addressing Service Attribute (see section 11.4) and the SD-WAN UNI IPv6 Connection Addressing Service Attribute (see section 11.5), there is no need to specify any additional routing information (static routing or dynamic routing protocols). In that case, the Subscriber can use a "default gateway", i.e., a default route towards the Service Provider's address specified in the SD-WAN UNI IPv4 Connection Addressing Service Attribute (see section 11.4) or the SD-WAN UNI IPv6 Connection Addressing Service Attribute (see section 11.5). As above, the Service Provider directs traffic towards the UNI that is destined for an IP address within the IP Prefix identified by the connection addressing attributes, and in this case the value of the SD-WAN UNI Routing Protocols Service Attribute can be an empty list.

Each of the routing protocols specified below has a parameter for setting the administrative distance. This is a numeric metric used to control which routes are selected when there are multiple routes for the same IP Prefix. A lower number indicates a more preferable route.

Setting the administrative distance can be useful when the same IP Prefix in the Subscriber Network is reachable via two or more UNIs, to allow the Subscriber to select which UNI the Service Provider should use to reach IP Addresses within the IP Prefix. Note that once a given UNI is selected, it remains the Service Provider's responsibility to direct IP Packets from the Ingress UNI to that Egress UNI over the UCSs in accordance with the Policy that is applied to those IP Packets. Setting the administrative distance only impacts which Egress UNI is used to reach the destination.

- [R140] IP Prefixes identified by the SD-WAN UNI IPv4 Connection Addressing Service Attribute (see section 11.4) and the SD-WAN UNI IPv6 Connection Addressing Service Attribute (see section 11.5) **MUST** be assigned an administrative distance of 0.
- [R141] For a given UNI, routes towards other UNIs **MUST** be assigned an administrative distance of 200.
- [R142] When selecting the best route for packet delivery as described in section 7.16, the Service Provider **MUST** select a route with the lowest administrative distance.

Note that the administrative distance values used in this document and specified in the value of the SD-WAN UNI Routing Protocols Service Attribute are only related to each other, to specify the relative preference of routes. They might or might not correspond with administrative distance values actually used in the Service Provider's devices to implement the behavior.

If two routes for an IP Prefix have the same administrative distance, then the Service Provider is free to choose which Egress UNI to use to deliver an IP Packet destined for that IP Prefix.

For BGP and OSPF, setting a different administrative distance for different IP Prefixes is not supported in this version of the standard.

11.6.1 Static

When an entry in the SD-WAN UNI Routing Protocols list is for *Static*, the IP Prefixes reachable in the Subscriber Network directly attached to this UNI are specified as additional parameters in the entry. These are known as Static IP Routes. For each Static IP Route, the following information is specified in addition to the IP Prefix:

- A next hop IP address in the Subscriber Network
- Administrative Distance, an integer greater than 0 and less than 256.

The Service Provider directs traffic destined for an address within any of the Static IP Routes towards the UNI, using the next hop address. The Subscriber routes traffic towards the UNI (e.g., by using a default or aggregate route).

The same IP Prefix can be specified more than once in the list of Static IP Routes if it has different forwarding information.

If a static prefix is specified with a next hop address that is not reachable over this UNI, the Static IP Route is considered inactive and hence is not used by the Service Provider for directing traffic. In particular, the Static IP Route is not used if the specified next hop can only be reached via a different UNI.

Note: if a Static IP Route is specified that matches the IP Prefix for the connection addresses for the UNI (see sections 11.4 and 11.5) the connected route is always preferred as it has administrative distance fixed to 0.

11.6.2 BGP

When an entry in the SD-WAN UNI Routing Protocols is for *BGP*, BGP as specified in RFC 4271 [20] is used across the UNI to exchange routing information. The Subscriber uses BGP to advertise IP Prefixes used within the Subscriber Network that are reachable over the UNI to the Service Provider, which consequently directs traffic destined for any IP address within those IP Prefixes towards the UNI corresponding to the next hop associated with the IP Prefix. The Service Provider uses BGP to advertise IP Prefixes that are reachable via other UNIs in the SWVC so that the Subscriber can direct traffic towards those IP Prefixes over the corresponding UNIs. The capabilities described in this section cover several common configurations including the five shown in Figure 19.

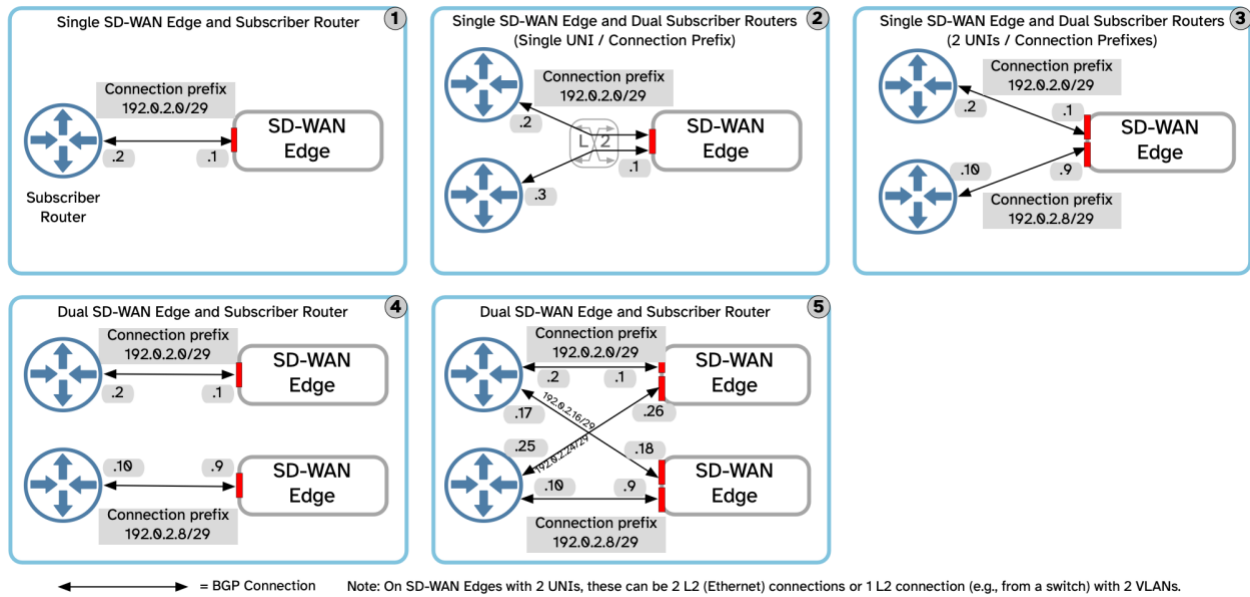


Figure 19 – BGP Configurations Examples

The Subscriber and the SD-WAN Edges use BGP to advertise and learn IP Prefixes that are reachable via UNIs in the SWVC, so that the Subscriber can direct traffic destined for those IP Prefixes towards the Service Provider, over the UNI corresponding to the next hop associated with the IP Prefix.

[R143] When an entry in the SD-WAN UNI Routing Protocols is specified for *BGP*, BGP as specified in RFC 4271 [20] **MUST** be used across the UNI to exchange routing information.

The additional parameters that need to be agreed on when BGP is used:

- The AS Number of the Subscriber
- The AS Number of the Service Provider
- Connection Address Family (*IPv4* or *IPv6*)
- A list of one or more Subscriber IP Addresses at the UNI that are:
 - In the address family identified by the Connection Address Family parameter
 - Within the Connection Primary IPv4 Prefix (see section 11.4) or a Connection IPv6 Prefix (see section 11.5) for the UNI, and
 - Not included in any of the Service Provider’s reserved addresses (which includes DHCP-assigned addresses)
- Authentication (*None* or MD5 plus a password)
- BGP Community List (may be empty) (see below)
- BGP Extended Community List (may be empty) (see below)
- Hold Time (time in seconds)
- Damping (*None* or a set of damping parameters)
- AS Override (Enabled or Disabled)
- Administrative Distance (integer greater than 0 and less than 256)

The Subscriber’s and Service Provider’s AS Numbers are used to establish BGP peerings.

One or more BGP sessions can be established between the Subscriber and the SD-WAN Edge. These sessions can be established using either IPv4 or IPv6, as specified by the Connection Address Family parameter. For each session a Subscriber IP Address within one of the connection prefixes (see sections 11.4 and 11.5) is agreed on. For IPv4 sessions, the session is established between this Subscriber address and the first Service Provider address in the primary subnet specified in the SD-WAN UNI IPv4 Connection Addressing Service Attribute (section 11.4). For IPv6 sessions, the session is established between the specified Subscriber address and the first Service Provider address in the first subnet specified in the SD-WAN UNI IPv6 Connection Addressing Service Attribute (section 11.5). These peering sessions are single hop eBGP connections. Note in this case that the same values for the parameters above apply to each of these BGP peering sessions.

- [R144] When an entry in the value of the SD-WAN UNI Routing Protocols Service Attribute is *BGP*, the Service Provider **MUST** support 4-octet AS Numbers as described in RFC 6793 [26].
- [R145] When an entry in the value of the SD-WAN UNI Routing Protocols Service Attribute is *BGP*, and the Connection Address Family parameter is IPv4, then the SD-WAN UNI IPV4 Connection Addressing Service Attribute (section 11.4) **MUST** be *Static*.
- [R146] When an entry in the value of the SD-WAN UNI Routing Protocols Service Attribute is *BGP*, and the Connection Address Family parameter is IPv6, then the SD-WAN UNI IPV6 Connection Addressing Service Attribute (Section 11.5) **MUST** be *Static*.
- [R147] When an entry in the value of the SD-WAN UNI Routing Protocols Service Attribute is *BGP*, if the Authentication parameter is MD5, authentication using MD5 **MUST** be used as described in RFC 4271 [20] using the specified password.

The Service Provider can configure BGP to wait passively for the Subscriber's devices to connect to it; this is helpful if it is not known whether the Subscriber devices are available yet. To ensure this works, the Subscriber has to use active mode.

- [R148] The Subscriber **MUST NOT** use passive TCP establishment for BGP sessions with the Service Provider.

The Service Provider, however, can use passive TCP establishment for BGP sessions with the Subscriber.

BGP Communities and Extended Communities allow additional metadata to be attached to route advertisements. Except in the case of the few standardized well-known values, this additional metadata has no intrinsic meaning. However, it is common for Service Providers to define a set of Communities or Extended Communities with associated semantics, that the Subscriber can attach to their route advertisements in order to affect how they are handled by the Service Provider.

- [R149] Each entry in the BGP Community List and BGP Extended Community List parameters **MUST** have an associated semantic that describes how the Service Provider will handle routes advertised with that value.

The Hold Time parameter indicates the agreed Hold Time used for the BGP sessions. The possible values are 0 or an integer in the range 3–65535, as defined in RFC 4271 [20].

The Service Provider can apply route flap damping to advertisements from the Subscriber, but in this case the parameters have to be agreed.

- [R150] When the Damping parameter is not *None*, the Service Provider **MUST** apply route flap damping as described in RFC 2439 [11].
- [R151] When the Damping parameter is not *None*, a single set of parameters as described in section 4.2 of RFC 2439 [11] **MUST** be agreed.
- [R152] When the Damping parameter is *None*, the Service Provider **MUST NOT** apply route flap damping.

In cases where the Subscriber uses the same AS number in different parts of the Subscriber Network, it is necessary to tweak the normal handling of AS Paths in routes advertised to the Subscriber at each UNI, to prevent the routes being discarded due to BGP's loop prevention mechanisms. Two mechanisms are commonly used for this:

- The Subscriber can configure their BGP routers to disable the loop prevention mechanism in the case where their own AS Number appears in the AS Path (this is commonly known as “Allow-AS-in”). In this case, the Service Provider does not need to be aware that this is being done and hence no parameters need to be agreed.
- The Service Provider can overwrite instances of the Subscriber's AS Number in the AS Path with their own AS Number, when advertising routes to the Subscriber (this is commonly known as “AS Override”). This needs to be explicitly agreed between the Service Provider and the Subscriber.

- [R153] When the AS Override parameter is Enabled, the Service Provider **MUST** overwrite all instances of the Subscriber's AS Number in the AS Path with their own AS Number, in routes advertised to the Subscriber.

The Administrative Distance is applied by the Service Provider to all IP Prefixes advertised by the Subscriber over the UNI using BGP.

11.6.3 OSPF

When an entry in the UNI Routing Protocols is for *OSPF*, then OSPF as specified in RFC 2328 [10] (for IPv4) and/or RFC 5340 [24] (for IPv6) is used across the UNI to exchange routing information. The Subscriber uses OSPF to advertise, to the Service Provider, IP Prefixes used within the Subscriber Network that are reachable via the UNI. The Service Provider consequently directs traffic destined for any IP address within those IP Prefixes towards the UNI over which the IP Prefixes were advertised. The Service Provider uses OSPF to advertise IP Prefixes that are

reachable via other UNIs in the SWVC, so that the Subscriber can direct traffic towards those IP Prefixes over the corresponding UNIs.

- [R154]** When an entry in the UNI Routing Protocols is specified for *OSPF*, then OSPF as specified in RFC 2328 [10] (for IPv4) and/or RFC 5340 [24] (for IPv6) **MUST** be used across the UNI to exchange routing information.

The additional parameters that need to be agreed when OSPF is used are:

- Area ID (0 – 4294967295, normally expressed as an IPv4 address)
- Area type (*Normal*, *Stub* or *NSSA*)
- Authentication Type (for IPv4: *None*, *Password*, or *Message Digest*)
- Hello Interval (0 – 65535, in seconds)
- Dead Interval (0 – 4294967295, in seconds)
- Retransmit Interval (integer greater than 0, in seconds)
- Administrative Distance (integer in the range 1–255)

The Area ID is a 32-bit number (typically written as an IPv4 address) that specifies the OSPF Area.

If the Area ID is 0 (0.0.0.0), the area is the OSPF Backbone area. This can be used at the UNI, for example, if the Subscriber wishes for the Service Provider to implement a “super backbone” configuration, which allows the remote networks to appear to be in the same OSPF Backbone Area (Area ID 0), preserving the Subscriber’s route types. If a “super-backbone” is not used, the Subscriber’s routes from the remote locations will be learned as external, which can affect the routing within the Subscriber Network.

The Area Type indicates the type of OSPF Area. An Area Type of Normal means the area is not a stub or NSSA (see RFC 3101 [15]).

The Authentication Type indicates the type of authentication used for IPv4 OSPFv2 adjacencies. Authentication for IPv6/OSPFv3 is beyond the scope of this document. The Hello Interval, Dead Interval, and Retransmit Interval specify the various timers that are used to create OSPF adjacencies.

The Administrative Distance is an integer greater than 0 and is applied by the Service Provider to all IP Prefixes advertised by the Subscriber over the UNI using OSPF.

Note: parameters, behavior and requirements relating to the use of OSPF Sham links, and further parameters relating to authentication are not addressed in this version of the standard.

11.7 SD-WAN UNI Availability Objective Service Attribute

UNI Availability is the proportion of time, during a given time period, T_k , as described in the SWVC Service Performance Objectives Reporting Periods Service Attribute (section 9.8), that the SD-WAN UNI is working from the perspective of the Subscriber, excluding any agreed exceptions, e.g., maintenance intervals. The value of this Service Attribute is *None* or a value, \overline{UA} , expressed as a percentage between 0% and 100% that is the objective for UNI Availability for each time period, T_k .

- [R155] If the value of the SD-WAN UNI Availability Objective Service Attribute is not *None* for any UNI, the SD-WAN Subscriber and Service Provider **MUST** agree on what constitutes a UNI outage, including determining when an outage starts and ends.
- [R156] If the value of the SD-WAN UNI Availability Objective Service Attribute is not *None* for any UNI, the SD-WAN Subscriber and Service Provider **MUST** agree on what constitutes an exception, including determining when an exception starts and ends.

Note that [R155] and [R156] require that the agreement on what constitutes a UNI outage and what constitutes an exception (such as a maintenance period) is the same for all UNIs (unless the value of the SD-WAN UNI Availability Objective Service Attribute is *None* for all UNIs in the SD-WAN).

- [R157] The UNI Availability during a time period T_k **MUST** be defined as follows:
- Let $O(T_k)$ be the total duration of outages for this UNI during the time period T_k .
 - Let $E(T_k)$ be the total duration of exceptions for this UNI, such as maintenance periods, during the time period T_k . (Note that $E(T_k)$ must be $\leq T$ (the duration of time period T_k)).
 - Then define the UNI Availability as follows:
 - If $T > E(T_k)$, $UA(T_k) = \frac{T - (E(T_k) + O(T_k))}{T - E(T_k)}$
 - If $T = E(T_k)$, $UA(T_k) = 1$

The definition of what constitutes a UNI outage is often (but does not have to be) based on the raising and resolution of customer complaints (trouble tickets) rather than the actual performance of data traffic across the UNI. This definition can be refined based on further commercial considerations, such as exceptions for events beyond the Service Provider's control. An exact definition is outside the scope of this document.

- [R158] If the value of the SD-WAN UNI Availability Objective Service Attribute is not *None*, the SD-WAN UNI Availability Objective for a given time interval, T_k , **MUST** be considered met for this UNI if and only if $UA(T_k) \geq \widehat{UA}$.

11.8 SD-WAN UNI Mean Time To Repair Objective Service Attribute

The SD-WAN UNI Mean Time To Repair (MTTR) Objective is the arithmetic mean, i.e., average, of the durations of all outages of the UNI that start in a given time period T_k (as described in the SWVC Service Performance Objectives Reporting Periods Service Attribute, section 9.8), excluding any agreed exceptions, e.g., maintenance intervals. This means that an outage that crosses the boundary of a reporting period (T_k) is counted for MTTR in the period in which it started.

The value of this Service Attribute is *None* or a value, \widehat{UR} , that is the objective for UNI Mean Time To Repair for each time period, T_k . \widehat{UR} is expressed as a time duration in minutes.

A period when a UNI is unavailable can be considered as a single outage or as multiple outages for the purpose of the MTTR calculation. For example, if outages are based on trouble tickets and a trouble ticket is raised for the UNI at $t=1$ and this is resolved at $t=4$, then this is a single outage with a duration of 3. If, however, a second trouble ticket is raised for the UNI at $t=3$ (while the first outage is active) and is resolved at $t=7$ there are at least two ways this can be computed. If this is considered as two outages, then the MTTR can be computed on the two independently (duration of 3 for the first outage and duration of 4 for the second outage). But if the issues identified in the two trouble tickets are treated as a single outage, then the MTTR is computed based on that single outage of duration 6. The result depends on how the Subscriber and the Service Provider agree that outages are reported and resolved.

- [R159] If the value of the SD-WAN UNI Mean Time To Repair Objective Service Attribute is not *None* for any UNI, the SD-WAN Subscriber and Service Provider **MUST** agree on the definition of an outage, including determining when an outage starts and ends.
- [R160] If the value of the SD-WAN UNI Mean Time To Repair Objective Service Attribute is not *None* for any UNI, the SD-WAN Subscriber and Service Provider **MUST** agree on what constitutes an exception, including determining when an exception starts and ends.

Note that [R159] and [R160] require that the agreement on what constitutes a UNI outage and what constitutes an exception (such as a maintenance period) is the same for all UNIs (unless the value of the SD-WAN UNI Mean Time To Repair Objective Service Attribute is *None* for all UNIs in the SD-WAN).

- [R161] If the value of the SD-WAN UNI Availability Objective Service Attribute (section 11.7) is not *None* for any UNI and the value of the SD-WAN UNI Mean Time To Repair Objective Service Attribute is not *None* for any UNI, then the agreement of what constitutes a UNI outage and an exception **MUST** be the same in both calculations.
- [R162] If the value of the SD-WAN UNI Mean Time To Repair Objective Service Attribute is not *None* for any UNI, then agreed exception periods **MUST NOT** overlap in time.

The effect of requirement [R162] is that any separate and overlapping exception periods are considered as one period that is the union of the separate ones. The requirement is to ensure that exception periods are subtracted from outages only once.

As a result of requirement [R159] each outage period, U , has an agreed-on start time, U_{start} , and end time, U_{end} . Since the exception periods are agreed on, each exception period, E , has an agreed-on start time, E_{start} , and end time, E_{end} .

- [R163] The UNI Mean Time To Repair during a time period T_k **MUST** be defined as follows:
 - Let $Od(U, T_k)$ be the duration of an outage U that starts in T_k ; that is, $Od(U, T_k) = U_{end} - U_{start}$, where U_{start} is during T_k .

- Consider any exception periods that overlap with outage U . Let $Ed(U, T_k)$ be the sum of the durations of the parts of these exception periods that overlap with outage U . If there are no exception periods that overlap with U , then $Ed(U, T_k)$ is 0.
- Let $Vd(U, T_k) = Od(U, T_k) - Ed(U, T_k)$ for each outage U that starts in T_k .

Then the Mean Time to Repair $UR(T_k)$ is the arithmetic mean of the values $Vd(U, T_k)$ for all outages U that start in T_k (i.e., U_{start} is during T_k) for which $Vd(U, T_k) > 0$. If there are no such outages, then $UR(T_k)$ is 0.

The definition of what constitutes a UNI outage is often (but does not have to be) based on the raising and resolution of customer complaints (trouble tickets) rather than the actual performance of data traffic across the UNI. This definition can be refined based on further commercial considerations, such as exceptions for events beyond the Service Provider's control. An exact definition is outside the scope of this document.

[R164] If the value of the SD-WAN UNI Mean Time To Repair Objective Service Attribute is not *None*, the SD-WAN UNI Mean Time To Repair Objective for a given time interval, T_k , **MUST** be considered met if and only if $UR(T_k) \leq \widehat{UR}$.

Figure 20 shows some of the common cases described by the definition of $UR(T_k)$:

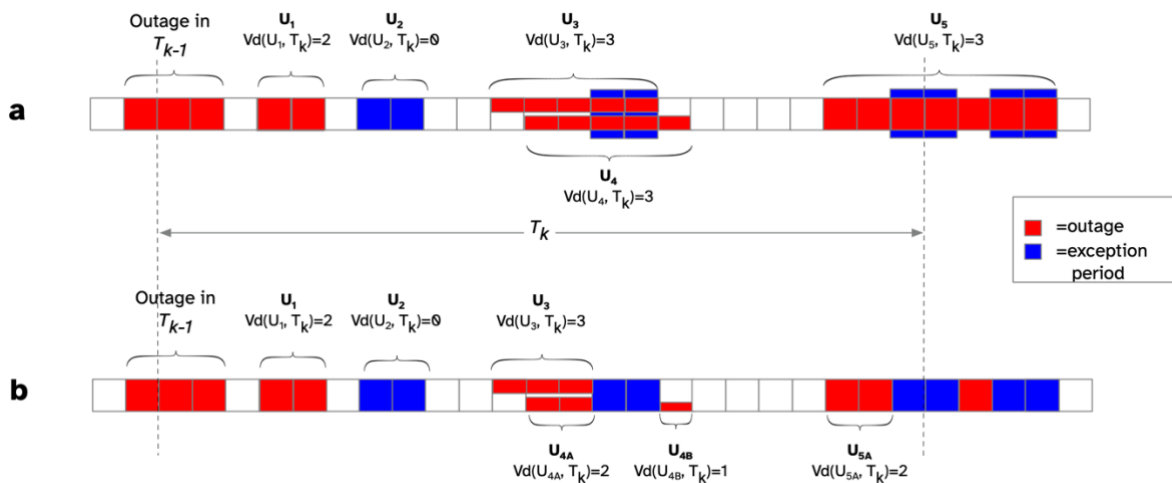


Figure 20 – Examples of Outages and Exception Periods

The MTTR calculation depends on an agreement between the Service Provider and Subscriber on how exception periods relate to outages. For example, exceptions could be defined independently of outages, such that exceptions and outages can overlap, and the exceptions are subtracted out in the MTTR calculation as noted in third bullet of [R163]. This is depicted in the top example (a) in Figure 20. Another possibility is that exception periods are considered to terminate any outages that precede them, such that exceptions and outages cannot overlap. If the UNI remains unavailable, a new outage could be considered to start at the end of the exception period. This is depicted in the bottom example (b). The examples in Figure 20 shows several outages and

exception periods. The time divisions shown in the diagram are to illustrate the duration of each outage and exception period—in practice outages and exception periods can start and end at any time.

The MTTR for the indicated period, T_k , can be determined as follows:

- The first period started during the period before T_k , so it is not included in the MTTR computation for T_k .
- The next period (U_1) is the simplest case, an outage with no overlapping exception periods. The duration of the outage is 2.
- Period U_2 is an isolated exception period, i.e., not adjacent to any outages, therefore it is not included in the MTTR calculation.
- Period U_3 is an outage that is adjacent to an exception period of duration 2. If the exception period is considered part of the outage, then the duration of the outage is 5 but $Vd(U_3, T_k) = 3$ after the exception period duration is subtracted out. This is shown in the (a). If the exception period is not considered part of the outage, the duration is 3 and $Vd(U_3, T_k) = 3$. This is shown in (b).
- Period U_4 is an outage that starts while outage U_3 is in progress (for example, it might have been reported on a separate trouble ticket). There is an initial component with duration 2 followed by an exception period of duration 2. When the exception period ends, the outage is still in progress. If the exception period is considered part of the outage, this is an outage with a duration of 5 but $Vd(U_4, T_k) = 3$ after the exception period duration is subtracted out. This is shown in (a). If the exception period is not considered part of the outage, then this could also be treated as two outages (shown as U_{4A} and U_{4B}), one with duration 2 and one with duration 1, likely resulting in a different value for MTTR. This is shown in (b).
- Period U_5 is the most complicated. The outage crosses the boundary into T_{k+1} , but since it started in T_k it is counted in the calculation of MTTR for T_k . It is overlapped by two exception periods, one that straddles the boundary between T_k and T_{k+1} and one that is entirely in T_{k+1} . If the exception periods are considered part of the outage, then this is a single outage with duration 7 but $Vd(U_5, T_k) = 3$ after the exception period duration is subtracted out. This is shown in (a). However, if the exception periods are not considered part of the outage, then the first part is an outage with duration $Vd(U_{5A}, T_k) = 2$ and the continuation of the outage after the exception period is in T_{k+1} so it is not counted in the MTTR for T_k . This is shown in (b).

Based on the example in Figure 20 and the subsequent bullet list, there are two computations for MTTR:

- If the exception periods are considered part of the outage (a), then $MTTR = UR(T_k) = \text{average}(Vd(U_1, T_k)=2, Vd(U_3, T_k)=3, Vd(U_4, T_k)=3, Vd(U_5, T_k)=3) = 11/4 = \mathbf{2.75}$, or
- If the exception periods are not considered part of the outage (b), then $MTTR = UR(T_k) = \text{average}(Vd(U_1, T_k)=2, Vd(U_3, T_k)=3, Vd(U_{4A}, T_k)=2, Vd(U_{4B}, T_k)=1, Vd(U_{5A}, T_k)=2) = 10/5 = \mathbf{2}$

12 Attributes for Underlay Connectivity Services

This section contains Service Attributes that apply to each of the Underlay Connectivity Services that compose the SD-WAN Service (SWVC). There is one instance of these attributes for each Underlay Connectivity Service underlying the SWVC.

Underlay Connectivity Services are network services independent of the SD-WAN Service and can have a large number of “characteristics” or “attributes” that define their configuration and behavior. If the UCS conforms to a MEF-defined Service, these characteristics and attributes are the MEF Service Attributes.

[R165] The SD-WAN Subscriber and the SD-WAN Service Provider **MUST** agree on the attributes and characteristics of the UCS that the SD-WAN Subscriber needs to communicate to the SD-WAN Service Provider.

[R166] The SD-WAN Subscriber **MUST** communicate the values of the attributes and characteristics of the UCS agreed per [R165] to the SD-WAN Service Provider, if the SD-WAN Service Provider is not already party to them.

Note that the SD-WAN Service Provider may have knowledge of the attributes and characteristics for the UCS either because they are also the UCS Service Provider, or because they purchased the UCS on behalf of the SD-WAN Subscriber. In these cases, there is no need for the SD-WAN Subscriber to communicate them.

The following sub-sections define additional UCS, UCS End Point, and UCS UNI Service Attributes that are specific to the use of the UCS for SD-WAN. The additional attributes include only those necessary to define UCS external interfaces and behavior to the extent necessary to implement SD-WAN Policies. Attributes that are internal to the UCS itself, such as those defining the associations between the UCS and UCS End Points or between the UCS End Point and UCS UNI, are part of the agreement between the UCS Subscriber and UCS Service Provider. These attributes are not included as Service Attributes for the SD-WAN Service and are communicated to the SD-WAN Service Provider per [R165] and [R166].

12.1 SD-WAN UCS Service Attributes

The SD-WAN UCS Service Attributes are summarized in the following table, and each is described in more detail in the subsequent sections.

Attribute Name	Summary Description	Possible Values
UCS Identifier	Identification of the Underlay Connectivity Service for management purposes.	Unique Identifier String for the SD-WAN Service
UCS Type	Indicates whether the UCS is a Public UCS (i.e., Internet Access Service) or a Private UCS.	<i>Public or Private</i>
UCS Billing Method	Indicates how the UCS is billed	<i>Flat-rate, Usage-based or Other</i>

Table 12 – Summary of SD-WAN UCS Service Attributes

12.1.1 UCS Identifier Service Attribute

The value of the UCS Identifier Service Attribute is a string that is used to allow the Subscriber and Service Provider to uniquely identify an Underlay Connectivity Service.

- [R167] The value of the UCS Identifier Service Attribute **MUST** be an Identifier String.
- [R168] The value of the UCS Identifier Service Attribute **MUST** be unique across all UCS Identifiers in the Service Provider Network.
- [R169] The value of the UCS Identifier Service Attribute **MUST** be the same as the value used to identify the UCS when communicating the Service Attributes or characteristics of the UCS per [R166].

12.1.2 UCS Type Service Attribute

The value of the UCS Type Service Attribute indicates whether the UCS is a Public UCS (i.e., an Internet Access Service) or a Private UCS. The possible values are *Public* and *Private*.

- [R170] If the Underlay Connectivity Service is an IP Service with characteristics consistent with the following MEF 61.1 [32] Service Attributes: IPVC Topology Service Attribute = *Cloud Access* and IPVC Cloud Service Attribute with Cloud Type parameter = *Internet Access*; then the value of the UCS Type Service Attribute **MUST** be *Public*.
- [R171] If the Underlay Connectivity Service is not an Internet Access Service as described in [R170], then the value of the UCS Type Service Attribute **MUST** be *Private*.

12.1.3 UCS Billing Method Service Attribute

The UCS Billing Method Service Attribute indicates how access to the Underlay Connectivity Service is billed. The allowed values are *Flat-rate*, *Usage-based*, and *Other*.

UCSs with the value *Other* for this Service Attribute are only used for Application Flows with policies where the BILLING-METHOD is *Either* or is not specified.

12.2 SD-WAN UCS UNI Service Attributes

Access to an Underlay Connectivity Service is provided at the SD-WAN Edge via the UCS UNI. MEF uses the term UNI consistently across all service standards to represent the demarcation point between the responsibility of the Subscriber and the responsibility of the Service Provider. Although not all Underlay Connectivity Services are MEF Services, the concept of this demarcation point is relevant to all carrier-based services. So, if the Underlay Connectivity Service is a MEF service the UCS UNI refers to the MEF-defined UNI for that service and if the Underlay Connectivity Service is not a MEF-defined service the UCS UNI refers, nonetheless, to the relevant demarcation point.

The SD-WAN UCS UNI Service Attributes are summarized in the following table, and each is described in more detail in the subsequent sections.

Attribute Name	Summary Description	Possible Values
UCS UNI Identifier	Identification of the Underlay Connectivity Service UNI for management purposes.	Identifier String

Table 13 – Summary of SD-WAN UCS UNI Service Attributes

12.2.1 UCS UNI Identifier Service Attribute

The value of the UCS UNI Identifier Service Attribute is a string that is used to allow the Subscriber and Service Provider to uniquely identify an Underlay Connectivity Service UNI.

- [R172] The value of the UCS UNI Identifier Service Attribute **MUST** be an Identifier String.
- [R173] The value of the UCS UNI Identifier Service Attribute **MUST** be unique across all UCS UNI Identifiers in the Service Provider Network.
- [R174] The value of the UCS UNI Identifier Service Attribute **MUST** be the same as the value used to identify the UCS UNI when communicating the Service Attributes or characteristics of the UCS per [R166].

12.3 SD-WAN UCS End Point Service Attributes

A UCS is connected to a UCS UNI by a UCS End Point. As discussed in section 7.7, some service architectures support multiple UCSs at the same UCS UNI. The UCS End Point captures Service Attributes at the UCS UNI whose value can differ between UCSs.

The SD-WAN UCS End Point Service Attributes are summarized in the following table and described in the subsequent sections.

Attribute Name	Summary Description	Possible Values
UCS End Point Identifier	Identification of the Underlay Connectivity Service End Point.	Identifier String
UCS End Point Backup	Indicates whether the UCS should be treated as a Primary underlay or a Backup underlay at this UCS End Point.	<i>Enabled or Disabled</i>
UCS End Point Breakout	Indicates whether the SD-WAN Service can “break out” to the underlying UCS Service from this End Point.	<i>Enabled or Disabled</i>

Table 14 – Summary of SD-WAN UCS End Point Service Attributes

12.3.1 UCS End Point Identifier Service Attribute

The value of the UCS End Point Identifier Service Attribute is a string that is used to allow the Subscriber and Service Provider to uniquely identify an Underlay Connectivity Service End Point.

- [R175] The value of the UCS End Point Identifier Service Attribute **MUST** be an Identifier String.
- [R176] The value of the UCS End Point Identifier Service Attribute **MUST** be unique across all UCS End Point Identifiers for SWVCs (SD-WAN Services) in the Service Provider network.
- [R177] The value of the UCS End Point Identifier Service Attribute **MUST** be the same as the value used to identify the UCS End Point when communicating the Service Attributes or characteristics of the UCS per [R166].

12.3.2 UCS End Point Backup Service Attribute

This Service Attribute allows the Service Provider and Subscriber to agree on whether the UCS (at this UCS End Point) performs the role of a Backup UCS for some Application Flows. The value of the UCS End Point Backup Service Attribute can be *Enabled* or *Disabled* indicating whether the UCS can be used for any Application Flow (*Disabled*) or only those flows whose Policy at this SD-WAN Edge indicates that they can use a Backup UCS. See section 9.12.2.7 for details about how the BACKUP Policy is used.

12.3.3 UCS End Point Breakout Service Attribute

Application Flows can be assigned Policies that allow them to be forwarded to IP destinations that are reachable via the UCS rather than via the Subscriber Network at an SD-WAN UNI. This operation is referred to as “breakout”. This Service Attribute allows the Subscriber and Service Provider to agree on the UCS End Points that allow breakout. Breakout can occur at all UCS End Points in the UCS or at specific UCS End Points. If breakout is disallowed at a UCS End Point, Application Flows can be forwarded over the SWVC (i.e., over TVCs in the SWVC) to another SD-WAN Edge connected to the UCS that allows breakout at the UCS End Point.

Although, in the general case, any UCS can support breakout, in this standard only Underlay Connectivity Services that are Internet Access Services (the value of the UCS Type Service Attribute is *Public*) support breakout (see Internet Breakout in section 7.14). Since all Internet Access UCSs are, in effect, connected, disallowing breakout at one UCS End Point implies that the Application Flows can be forwarded to another SD-WAN Edge with an Internet Access UCS with an End Point that allows breakout.

The value of the UCS End Point Breakout Service Attribute can be *Enabled* or *Disabled* and indicates whether breakout is allowed at the UCS End Point.

- [R178] If the value of the UCS End Point Breakout Service Attribute is *Disabled*, IP Packets **MUST NOT** be forwarded over the UCS EP to destinations reachable via the UCS, other than to other SD-WAN Edges in the SD-WAN Service via a TVC.

Defining this Service Attribute as a generic “breakout” Service Attribute rather than specifically an Internet Breakout Service Attribute, provides flexibility for the Service Provider to define additional breakout-related Policy Criteria and capabilities.

13 Performance Metrics

The PERFORMANCE Policy Criterion (see section 9.12.2.8) provides the means for Policies to specify real-time performance requirements for Application Flows. The SD-WAN Service attempts to forward IP Packets over Paths that meet or exceed the performance goals specified in the PERFORMANCE Policy Criterion, consistent with the requirements and restrictions posed by other Policy Criteria. The Performance Criterion supports three different Performance Metrics: One-way Mean Packet Delay, One-way Mean Packet Delay Variation, and One-way Packet Loss Ratio.

Section 13.1 defines Qualified Packets that are the packets for which the Performance metrics apply. Section 13.2 defines One-Way Packet Delay, on which two of the three metrics are based. The subsequent three subsections define the three Performance Metrics used in this document.

13.1 Qualified Packets

The Performance Metrics defined in the sections below apply to Qualified Packets. A Qualified Packet between two UNIs, referred to as x and y , across a Path, referred to as p , is any unicast IP Packet that satisfies the following conditions:

- The IP Packet ingresses at UNI x
- The IP Packet is destined to egress at UNI y
- The IP Packet is not discarded per requirements [R23], [R38], [R64], [R83], [R85], [R96], [R112], [R113], [D8], [D1]
- The Path p is chosen to carry the IP Packet
- The IP Packet is not fragmented
- The IP Packet does not incur significant intentional delay as a result of applying a Security Policy agreed to between the Service Provider and the Subscriber.

The inclusion of [R38] in the third bullet above covers all cases in which policy or routing prevents the forwarding of a packet, including discarding packets to enforce virtual topologies, allowed destination zones, or security policies. Such packets are therefore never considered to be Qualified Packets.

The final bullet in the list above is focused on delays as a result of Security Functions (such as those described in MEF 88 [38]) that require extended analysis (e.g., quarantining or malware detection and removal). Small increases in delay due to processing are expected to be factored into the agreed-on performance goals for the Service.

13.2 One-Way Packet Delay

The One-way Packet Delay for a Qualified Packet that is sent from UNI x to UNI y across Path p is defined as the time elapsed from the reception of the first bit of the packet at the Ingress UNI until the transmission of the last bit of the first corresponding Packet at the Egress UNI. If the packet is duplicated in transit (either erroneously or intentionally, e.g., packet loss remediation), the delay is based on the first copy that is delivered.

Note that this definition of One-way Packet Delay for a packet includes the delays encountered as a result of transmission across the Ingress and Egress UNIs as well as that introduced by the network that connects them.

13.3 One-Way Mean Packet Delay Performance Metric

[R179] If a PERFORMANCE Policy Criterion includes a reference to One-Way Mean Packet Delay in the *primary* or *secondary* element, it **MUST** be defined as follows for each Path p between UNIs x and y :

Let $\Delta = \{\delta_1, \delta_2, \delta_3, \dots, \delta_n\}$ represent the One-Way Packet Delays of the n Qualified Packets sent from UNI x to UNI y across Path p during a time interval whose duration is the value of the *evalinterval* element of the SWVC Performance Time Intervals Service Attribute. Then the One-Way Mean Packet Delay for p over that interval is the arithmetic mean of the values $\delta_1 \dots \delta_n$. If $n=0$ during the time interval, the One-Way Mean Packet Delay for that time interval is zero.

13.4 One-Way Mean Packet Delay Variation Performance Metric

[R180] If a PERFORMANCE Policy Criterion includes a reference to One-Way Mean Packet Delay Variation in the *primary* or *secondary* element, it **MUST** be defined as follows for each Path p between UNIs x and y :

Let $\Delta = \{\delta_1, \delta_2, \delta_3, \dots, \delta_n\}$ represent the One-Way Packet Delays of the n Qualified Packets sent from UNI x to UNI y across Path p during a time interval whose duration is the value of the *evalinterval* element of the SWVC Performance Time Intervals Service Attribute. Let Δ' be the set of all pairs of elements $\{\delta_r, \delta_s\}$ in Δ such that $s > r$ and the difference in the arrival time at the Ingress UNI of packets s and r equals the value of the *arrivalinterval* element in the SWVC Performance Time Intervals Service Attribute. If Δ' is *null*, then the One-Way Mean Packet Delay Variation for the time interval is zero. Otherwise, let v_{rs} be the absolute value of the difference in One-Way Packet Delay for each pair, $\{\delta_r, \delta_s\}$ in Δ' , i.e., $v_{rs} = |\delta_r - \delta_s|$. Then the One-Way Mean Packet Delay Variation for p over that interval is the arithmetic mean of the values v_{rs} for each element in Δ' .

13.5 One-Way Packet Loss Ratio Performance Metric

[R181] If a PERFORMANCE Policy Criterion includes a reference to One-Way Packet Loss Ratio in the *primary* or *secondary* element, it **MUST** be defined as follows for each Path p between UNIs x and y :

Let s represent the total number of Qualified Packets sent from UNI x to UNI y across Path p during a time interval whose duration is the value of the *evalinterval* element of the SWVC Performance Time Intervals Service Attribute. Let r represent the total number of unique (not duplicate) Qualified Packets received from UNI x at UNI y on p that were sent during the same

period. Then the One-Way Packet Loss Ratio over that interval for p 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 $(s-r)/s$

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

14 References

- [1] IANA, *Protocol Numbers*,
<https://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml>
- [2] IANA, *Service Name and Transport Protocol Port Number Registry*,
<https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml>
- [3] IEEE Std 802.1Q-2018, *IEEE Standard for Local and Metropolitan Area Networks – Bridges and Bridged Networks*, May 2018
- [4] IEEE Std 802.3-2018, *IEEE Standard for Ethernet*, August 2018
- [5] IETF BCP78 (RFC5378), *Rights Contributors Provide to the IETF Trust*, by S. Bradner and J Contreras, November 2008. Copyright © 2008 IETF Trust and the persons identified as authors. All rights reserved.
- [6] IETF RFC 791, *Internet Protocol*, September 1981
- [7] IETF RFC 2119, *Key words for use in RFCs to Indicate Requirement Levels*, by Scott Bradner, March 1997
- [8] IETF RFC 2131, *Dynamic Host Configuration Protocol*, March 1997
- [9] IETF RFC 2132, *DHCP Options and BOOTP Vendor Extensions*, March 1997
- [10] IETF RFC 2328, *OSPF Version 2*, by John Moy, April 1998. Copyright © The Internet Society (1998). All Rights Reserved.
- [11] IETF RFC 2439, *BGP Route Flap Damping*, by Curtis Villamizar and Ravi Chandra and Dr. Ramesh Govindan, November 1998. Copyright © The Internet Society (1998). All Rights Reserved.
- [12] IETF RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*, Fred Baker and David L. Black and Kathleen Nichols and Steven L. Blake, December 1998. Copyright © The Internet Society (1998). All Rights Reserved.
- [13] IETF RFC 2764, *A Framework for IP Based Virtual Private Networks*, by Andrew G. Malis and Dr. Arthur Y. Lin and Dr. Juha Heinanen and Bryan Gleeson and Dr. Grenville Armitage, February 2000. Copyright © The Internet Society (2000). All Rights Reserved.
- [14] IETF RFC 3022, *Traditional IP Network Address Translator (Traditional NAT)*, by P. Srisuresh and K. Egevang, January 2001. Copyright © The Internet Society (2001). All Rights Reserved.
- [15] IETF RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*, by Dr. Patrick W. Murphy, January 2003. Copyright © The Internet Society (2003). All Rights Reserved.

- [16] IETF RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, by Sally Floyd and Dr. K. K. Ramakrishnan and David L. Black, September 2001. Copyright © The Internet Society (2001). All Rights Reserved.
- [17] IETF RFC 3260, *New Terminology and Clarifications for Diffserv*, by Daniel B. Grossman, April 2002. Copyright © The Internet Society (2002). All Rights Reserved.
- [18] IETF RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*, by Charles E. Perkins and Bernie Volz and Ted Lemon and Michael Carney and Jim Bound, July 2003. Copyright © The Internet Society (2003). All Rights Reserved.
- [19] IETF RFC 3550, *RTP: A Transport Protocol for Real Time Applications*, by Henning Schulzrinne and Stephen L. Casner and Ron Frederick and Van Jacobson, July 2003. Copyright © The Internet Society (2003). All Rights Reserved.
- [20] IETF RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*, by Yakov Rekhter and Susan Hares and Tony Li, January 2006. Copyright © The Internet Society (2006). All Rights Reserved.
- [21] IETF RFC 4342, *Profile for Datagram Congestion Control Protocol (DCCP)*, by Jitendra Padhye and Sally Floyd and Eddie Kohler, March 2006. Copyright © The Internet Society (2006). All Rights Reserved.
- [22] IETF RFC 4862, *IPv6 Stateless Address Autoconfiguration*, by Dr. Thomas Narten and Tatsuya Jinmei and Dr. Susan Thomson, September 2007. Copyright © The IETF Trust (2007).¹⁵
- [23] IETF RFC 4960, *Stream Control Transmission Protocol*, by Randall R. Stewart, September 2007. Copyright © The IETF Trust (2007).¹⁵
- [24] IETF RFC 5340, *OSPF for IPv6*, by Dennis Ferguson and Acee Lindem and John Moy, July 2008. Copyright © The IETF Trust (2008).¹⁵
- [25] IETF RFC 6274, *Security Assessment of the Internet Protocol Version 4*, July 2011. Copyright © 2011 IETF Trust and the persons identified as the document authors. All rights reserved.
- [26] IETF RFC 6793, *BGP Support for Four-Octet Autonomous System (AS) Number Space*, by Quaizar Vohra and Enke Chen, December 2012. Copyright © 2012 IETF Trust and the persons identified as the document authors. All rights reserved.
- [27] IETF RFC 8174, *Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words*, by Barry Leiba, May 2017. Copyright © 2017 IETF Trust and the persons identified as the document authors. All Rights Reserved.

¹⁵ The documents that include a reference to this footnote are subject to the rights, licenses and restrictions contained in [BCP 78](#) [5], and except as set forth therein, the authors retain all their rights.

- [28] IETF RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*, by Dr. Steve E. Deering and Bob Hinden, July 2017. Copyright © 2017 IETF Trust and the persons identified as the document authors. All Rights Reserved.
- [29] MEF 6.3, *Subscriber Ethernet Services Definitions*, November 2019
- [30] MEF 10.4, *Subscriber Ethernet Service Attributes*, Phase 4, December 2018
- [31] MEF 23.2, *Carrier Ethernet Class of Service – Phase 3*, August 2016
- [32] MEF 61.1, *IP Service Attributes*, January 2019
- [33] MEF 61.1.1, *Amendment to MEF 61.1: UNI Access Link Trunks, IP Addresses, and Mean Time to Repair Performance Metric*, July 2022
- [34] MEF 63, *Subscriber Layer 1 Service Attributes*, August 2018
- [35] MEF 69.1, *Subscriber IP Service Definitions*, February 2022
- [36] MEF 70.1, *SD-WAN Service Attributes and Service Framework*, November 2021
- [37] MEF 74, *Commercial Affecting Attributes Technical Standard*, December 2018
- [38] MEF 88, *Application Flow Security for SD-WAN Services*, November 2021
- [39] MEF White Paper, *Understanding SD-WAN Managed Services*, July 2017

Appendix A Processing Application Flows (Informative)

This appendix describes the external behavior expected at the SD-WAN Edge in two ways. The first section describes a hypothetical processing architecture. It does not imply that any implementation be architected or implemented based on this model. The second part describes the behavior from an IP Packet and Service Attribute point of view.

A.1 Process View

There are several steps in determining how to forward IP Packets as shown in Figure 21.¹⁶

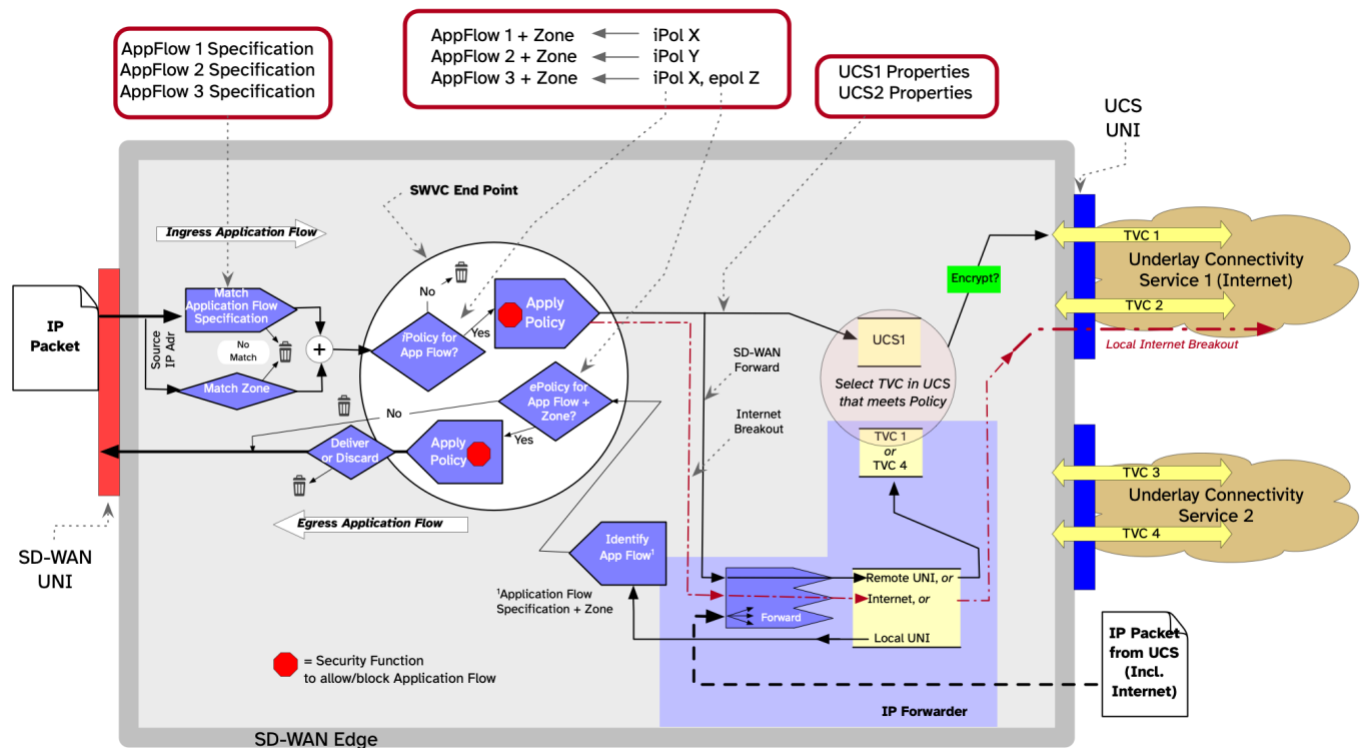


Figure 21 – Application Flows, Policies, and Forwarding

Each Ingress IP Packet is inspected to associate it with an Application Flow. This requires identifying which Application Flow Specification it matches. The source IP Address of the IP Packet is used to select the Zone. If the packet does not match any of the defined Application Flow Specifications or if the source address doesn't match a Zone, it is discarded. The Policy assigned to the Application Flow is then applied to the IP Packet (if no Policy is assigned, the IP Packet is discarded). Based on the Policy and the properties of the Underlay Connectivity Services and TVCs, a list of TVCs that can carry the IP Packet (in the diagram, TVC1 in UCS1) is determined.

¹⁶ This example assumes that the Service Provider has created all of the necessary TVCs to meet the Policy and connectivity requirements of the SD-WAN Service, and that the necessary IP routing information has been agreed to and configured.

The IP Forwarder determines where to forward this IP Packet. There are four possibilities:

- There is no route to the destination and the IP Packet is discarded
- The destination is at a remote UNI, i.e., the IP Packet must be forwarded across the SD-WAN Service. This results in a list of TVCs.
- The destination is on the Internet, i.e., the IP Packet must be forwarded to an Internet UCS if Policy allows Internet Breakout. The Internet Breakout can be local if an Internet UCS is available at the SD-WAN Edge or the IP Packet can be forwarded over the SD-WAN Service (i.e., a TVC) to another SD-WAN Edge for Internet Breakout.
- The destination is a UNI at this SD-WAN Edge and the IP Packet is forwarded to a local UNI if the Application Flow Policy allows.

If the destination is at a remote UNI, the forwarding process results in a list of TVCs that can reach the destination (in the diagram, TVC1 and TVC4). The intersection of this list and the TVCs that meet the Policy determines the TVC (or TVCs) that can carry the IP Packet (in this case, TVC1). Of course, it is possible for the intersection to include more than one TVC, in which case the Service Provider can choose any of them—the choice could be based on load balancing, performance optimization, or other criteria.

Packets arriving from the Underlay Connectivity Services go through a similar process. The IP Forwarder determines the destination (which can be the same as in the list above). If the IP Packet is destined for a remote UNI or the Internet (i.e., this is just an intermediate hop) then it is forwarded as appropriate. If the IP Packet is destined for a local UNI, the Application Flow Policy can be applied to the IP Packet to determine if it can be delivered to the UNI.

- If no Egress Policy is associated with the Application Flow, the IP Packet is forwarded to the UNI (note that this is the opposite of what happens for Ingress)
- If an Egress Policy is associated with the Application Flow, the IP Packet is forwarded or blocked based on that Policy.

An Egress Policy is determined based on the Application Flow to which the IP Packet is assigned at the Egress UNI. As noted in section 10.4.2, packets arriving from the Internet are assigned to the reserved Zone named *Internet*.

Figure 21 is a conceptual representation of the SD-WAN Edge functionality which is attempting to describe the externally visible behavior. In a given implementation the various processes could run in parallel, or they could be sequential or a mix. The description is implicitly assuming forwarding based on standard IP routing, but other IP forwarding approaches are possible, such as Policy-Based Routing, in which case this part of the process might be included in the “Apply Policy” process. The details of the implementation are beyond the scope of this document. The relevant point is that an IP Packet arrives, and one of the following actions can be taken:

1. It is forwarded over a TVC to another SD-WAN Edge
2. It is forwarded to a local UNI
3. It is forwarded to the Internet using Internet Breakout
4. It is discarded

A.2 Ingress Packet Flow View

The following diagram depicts the sequence of Service Attribute definitions that operate on an Ingress IP Packet.

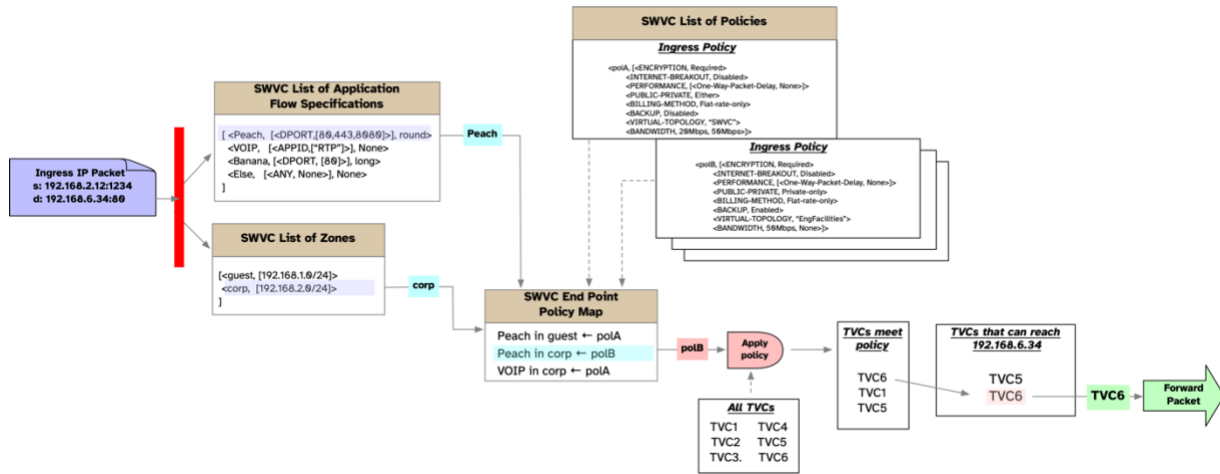


Figure 22 – Ingress Packet Flow

An Ingress IP Packet is assigned to an Application Flow based on the Application Flow Specification and the Zone. These two attributes of the IP Packet result in the selection of a Policy. In this case the IP Packet is associated with Application Flow *Peach* in Zone *corp*. This combination selects Policy *polB*. Policy *polB* can be met by TVCs 6, 1, and 5, ordered by the performance requirements specified by the PERFORMANCE Policy Criterion. TVCs 5 and 6 can reach the destination and therefore TVC6 is used because it has better performance.

Appendix B SD-WAN Use Cases (Informative)

This section provides several SD-WAN use cases to assist in putting the normative sections of this document into context. These diagrams and text have been derived from the MEF white paper, *Understanding SD-WAN Managed Services* [39].

B.1 Hybrid WAN




Figure 23 – Use Case: Hybrid WAN

Figure 23 illustrates a use case for an SD-WAN Service operating over three Underlay Connectivity Services, an Internet UCS (actually this is two UCS since each Internet Access Service is a different UCS) and an IP VPN UCS (e.g., implemented over MPLS) between the two sites. This hybrid UCS use case enables the Subscriber to use the multiple UCSs to achieve higher resiliency.

This is, perhaps, one of the most popular use cases because many enterprise Subscribers have both Internet and IP VPN UCSs to interconnect their sites, so the SD-WAN managed service enables them to take advantage of the benefits that SD-WAN provides over multiple UCSs.

B.2 Dual Internet WAN




Figure 24 – Use Case: Dual Internet WAN

Figure 24 illustrates a use case for an SD-WAN service operating over multiple Internet Service Providers (ISPs) to achieve resiliency using multiple Internet Underlay Connectivity Services plus different ISPs. The ISPs’ Internet connections could be DSL, Cable Internet, a dedicated Internet Access (DIA) service, or a combination of these.

Because the ISPs may not be the SD-WAN Service Provider, this use case could be applied to a larger SD-WAN managed service deployment where both sites are off-net and can only be reached via the Internet WAN.

Appendix C PERFORMANCE Policy Criterion Examples (Informative)

This appendix provides a sequence of examples on how the PERFORMANCE Policy Criterion (see section 9.12.2.8) can affect the choice of forwarding path and the invocation of performance remediation.

The examples focus on a single Application Flow at a UNI that is assigned a Policy with the following PERFORMANCE Policy Criterion:

(PERFORMANCE (“One-Way Mean Packet Delay”, 10ms, 12ms, 9ms), (“One-Way Packet Loss Ratio, 5%))

There are three available paths from the Ingress UNI to the Egress UNI for the Application Flow. A performance measurement at an arbitrary point in time, tI , has the following results:

Path	One-Way Mean Packet Delay	One-Way Packet Loss Ratio	Comment
A	8ms	4%	Best choice
B	11ms	1%	Delay exceeds <i>threshold</i>
C	10ms	3%	Not as good as A

Table 15 – Example measured values of Performance Metrics at time tI

Paths A and C meet the threshold value for both the primary Performance Metric (delay) and the secondary Performance Metric (Packet Loss Ratio). Path B exceeds the *threshold* for Packet Delay. Since A has a lower Mean Delay, it is the best choice although C might be chosen for other reasons. These choices (and similarly for the subsequent cases) are driven by [R69], [R70], and [D5].

At the next performance measurement at $tI+\Delta t$, the performance of the paths has changed as follows:

Path	One-Way Mean Packet Delay	One-Way Packet Loss Ratio	Comment
A	9ms	4%	Not as good as C
B	5ms	6%	Packet Loss Ratio exceeds <i>threshold</i>
C	8ms	1%	Best choice

Table 16 – Example measured values of Performance Metrics at time $tI+\Delta t$

Path B still does not meet the Policy, but in this case, it is because of Packet Loss not Delay. Paths A and C both meet the Policy with C being a bit better, but the *ceiling* value for Packet Delay is 12ms and therefore the Service Provider is biased towards keeping the Application Flow on Path A (and the subsequent examples assume that the Application Flow remains on Path A). However, the *remediation* value is 9ms which is the current delay, so the Service Provider should invoke performance remediation for Mean Packet Delay if it does not change to Path C.

At $tI+2\Delta t$, the performance of the paths has changed, again, as follows:

Path	One-Way Mean Packet Delay	One-Way Packet Loss Ratio	Comment
A	11ms	4%	Mean Packet Delay exceeds <i>threshold</i>
B	11ms	6%	Mean Packet Delay and Packet Loss Ratio exceeds <i>threshold</i>
C	9ms	5%	Only path that meets <i>threshold</i> values

Table 17 – Example measured values of Performance Metrics at time $tI+2\Delta t$

This is a relatively small change from the previous measurement period; the delay has degraded on all three Paths. Path C is now the only Path that meets the *threshold* values for both Performance Metrics. The Mean Packet Delay on Path A has increased in spite of remediation, but it is still below the *ceiling*, so the Service Provider can keep the Application Flow on Path A with remediation continuing. If the Application Flow were to be moved to Path C, remediation should be initiated there since 9ms is the trigger for that.

The final measurement at $tI+3\Delta t$ results in the following:

Path	One-Way Mean Packet Delay	One-Way Packet Loss Ratio	Comment
A	13ms	4%	Mean Packet Delay exceeds <i>ceiling</i>
B	5ms	6%	Packet Loss Ratio exceeds <i>threshold</i>
C	9ms	4%	Only path that meets <i>threshold</i> values

Table 18 – Example measured values of Performance Metrics at time $tI+3\Delta t$

Path A now has a Mean Packet Delay of 13ms. This exceeds the *ceiling* value for the Performance Metric and therefore the Service Provider is required to move the Application Flow to another Path that meets the Policy, if available. The delay on Path B has dropped to 5ms, but since the Packet Loss on Path B exceeds the *threshold* value, it is put at the end of the list of candidates for an alternative path per [R69]. Thankfully, Path C meets the Policy, and the Application Flow is moved to Path C. Since the Mean Packet Delay on Path C is at the *remediation* value, performance remediation should be initiated on this Path.

Appendix D Major changes from MEF 70.1 to MEF 70.2 (Informative)

The following list represents the major changes in MEF 70.2 from the previous version, MEF 70.1:

- In the UCS Billing Method Service Attribute, *Other* was removed as a possible value.
- Removed the SWVC Service Uptime Service Attribute.
- Removed the requirement that every Policy include the same set of Policy Criteria and added/modified text in each Policy Criterion to define behavior when a Policy does not include the Policy Criterion.
- New SD-WAN UNI Service Attributes: SD-WAN UNI Availability Objective and SD-WAN UNI Mean Time To Repair Objective
- New SWVC End Point Network Address Translation Service Attribute
- Enhanced the PERFORMANCE Ingress Policy Criterion to include values for *ceiling* and *remediation* parameters.
- New SWVC Service Attribute: SWVC Service Objectives Reporting Periods Service Attribute
- Introduced Rate Limiter as a new “key concept” in section 7.15.
- Added new SWVC Service Attribute: SWVC List of Rate Limiters and modified the BANDWIDTH Policy Criterion to allow specification of a named Rate Limiter.
- Changed how Application Flows can share bandwidth. They use named Rate Limiters instead of belonging to the same Application Flow Specification Group.
- Added new Appendix with examples of the PERFORMANCE Policy Criterion (Appendix C).
- Included normative text in section 12 about UCS Service Attributes of which the SD-WAN Service Provider must be aware.
- Added new SWVC Service Attribute: SWVC List of Application Flow Business Importance Levels.
- Added new Ingress Policy Criterion: BUSINESS-IMPORTANCE.
- In MEF 70.1 the value of the SWVC List of Security Policies Service Attribute (section 9.10) could be *None*, in MEF 70.2 that value is represented by an empty list.
- Changed the effect of INTERNET-BREAKOUT=Enabled. IP Packets in the Application Flow can be forwarded to Internet destinations or to UNI destinations based on their IP Address. In MEF 70.1 they could only be forwarded to Internet destinations.
- Changed the value of ALLOWED-DESTINATION-ZONES to allow an empty list. This can be used with INTERNET-BREAKOUT=Enabled to achieve behavior compatible with MEF 70.1 (see previous bullet).
- The BLOCK-SOURCE Egress Policy Criterion changed to EGRESS-BLOCK with values *Allow* and *Discard*.

Appendix E Acknowledgements (Informative)

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

- David **BALL**
- Bill **BJORKMAN**
- Neil **DANILOWICZ**
- Charles **ECKEL**
- Philippe **LOCHON**
- Basil **NAJEM**
- Larry **SAMBERG**
- John **SIWKO**