

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 23 January 2026

L. Dunbar
Futurewei
S. Hares
Huawei
K. Majumdar
Oracle
R. Raszuk
Arrcus
V. Kasiviswanathan
Arista
22 July 2025

BGP UPDATE for SD-WAN Edge Discovery
draft-ietf-idr-sdwan-edge-discovery-25

Abstract

The document describes the BGP mechanisms for SD-WAN (Software Defined Wide Area Network) edge node attribute discovery. These mechanisms include a new tunnel type and sub-TLVs for the BGP Tunnel-Encapsulation Attribute [RFC9012] and set of NLRI (network layer reachability information) for SD-WAN underlay information.

In the context of this document, BGP Route Reflector (RR) is the component of the SD-WAN Controller that receives the BGP UPDATE from SD-WAN edges and in turn propagates the information to the intended peers that are authorized to communicate via the SD-WAN overlay network.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 23 January 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	4
1.1. Secure L3VPN Services over SD-WAN	4
1.2. SD-WAN Secure Links	5
1.3. Conventions used in this document	5
2. Framework of BGP SD-WAN Edge Discovery	6
2.1. Supported Client Routes	7
2.2. Topologies of SD-WAN Edge Discovery	7
2.2.1. SD-WAN Secure L3VPN Example Topology	7
2.2.2. SD-WAN Secure Links Example Topology	11
2.2.3. RR connections to BGP Peers in SD-WAN	12
2.3. The Objectives of SD-WAN Edge Discovery	13
2.4. Examples of SD-WAN Edge Discovery	14
2.5. Comparing SD-WAN Secure L3VPN with Pure IPsec VPN	15
2.6. SD-WAN Segmentation, Virtual Topology and Client VPN	17
3. BGP SD-WAN Mechanisms	19
3.1. SD-WAN Hybrid Tunnel Encoding	19
3.2. SD-WAN Underlay UPDATE	23
3.2.1. The NLRI for SD-WAN Underlay Tunnel Update	23
3.2.2. Validation of SD-WAN NLRI	25
3.2.3. BGP Path Attributes attached to SD-WAN NLRI	26
3.3. IPsec-SA Property Sub-TLVs	26
3.3.1. IPsec-SA-ID Sub-TLV	26
3.3.2. IPsec-SA Rekey Counter Sub-TLV	28
3.3.3. IPsec Public Key Sub-TLV	30
3.3.4. IPsec-SA Proposal Sub-TLV	31

3.3.5.	Simplified IPsec-SA Sub-TLV	33
3.3.6.	Extended Port Attribute Sub-TLV	35
3.4.	Procedure for Client Routes with Hybrid SD-WAN Tunnel . .	41
3.4.1.	SD-WAN Hybrid Tunnel Type in Encapsulation Extended Community	42
3.4.2.	SD-WAN Hybrid Type in Tunnel Attributes via TEA . . .	43
3.4.3.	Client Routes Carried Over Multiple Hybrid SD-WAN Tunnels	44
3.4.4.	SD-WAN VPN ID in Control Plane	44
3.4.5.	SD-WAN VPN ID in Data Plane	45
3.5.	Procedure for Underlay Routes with Hybrid SD-WAN Tunnel TLV	45
3.5.1.	Hybrid SD-WAN NLRI without Encapsulation Extended Community	45
3.5.2.	Underlay Route with a TEA	45
3.5.3.	Underlay Routes with Port-Local-ID of Zero	47
3.5.4.	Multiple Tunnels attached to One Underlay Route . . .	47
3.6.	Error handling	48
3.6.1.	Error handling for the Tunnel Encapsulation Signaling	48
3.6.2.	Error Handling for NLRI	49
3.6.3.	SDWAN NLRI and Tunnel Encapsulation Attribute	50
4.	Operational Consistency and Tunnel Validation	50
4.1.	Detecting Misaligned Tunnels	50
4.2.	IPsec Attributes Mismatch	51
4.2.1.	Example creation of IPsec-SA over SD-WAN Hybrid tunnel	52
5.	Manageability Considerations	53
6.	Security Considerations	53
7.	IANA Considerations	54
7.1.	SD-WAN Overlay SAFI	54
7.2.	Tunnel Encapsulation Attribute Tunnel Type	54
7.3.	Tunnel Encapsulation Attribute Sub-TLV Types	55
7.4.	SD-WAN Edge Discovery NLRI Route Types	55
7.5.	SD-WAN Extended Port Encapsulation Types	55
7.6.	SD-WAN Extended Port Connection Types	56
7.7.	SD-WAN Extended Port Physical Port Types	56
8.	References	56
8.1.	Normative References	56
8.2.	Informative References	58
Appendix A.	Acknowledgments	59
Contributors	60
Authors' Addresses	60

1. Introduction

BGP [RFC4271] can serve as the control plane for a SD-WAN to support Secure VPNs or simply to support a set of Secure Links in a network. This section provides an overview of (1) L3VPN services supported by SD-WAN overlays and (2) SD-WAN Secure Links as an alternative tunneling mechanism for simpler deployments. Section 3 describes the framework of BGP SD-WAN edge discovery in terms of NLRIs supported, example topologies, and objectives for the BGP mechanisms. Section 4 describes the BGP mechanisms and corresponding error handling.

The BGP mechanisms supporting SD-WAN require that the RR establish a secure transport connection with each SD-WAN edge operating under the same BGP control plane instance. The establishment of a secure transport connection between each BGP Peer and the RR is outside the scope of this specification.

This document describes the BGP mechanisms for SD-WAN edge nodes to establish a BGP peering with other SD-WAN edge nodes, and pass information in order to establish and update SD-WAN overlay tunnels, as described in [Net2Cloud]. These mechanisms include a new tunnel type and sub-TLVs for the BGP Tunnel-Encapsulation Attribute [RFC9012] and a set of NLRIs for SD-WAN underlay information.

1.1. Secure L3VPN Services over SD-WAN

A SD-WAN network defined in [MEF70.1] and [MEF70.2], refers to a policy-driven network over multiple heterogeneous underlay networks to get better WAN bandwidth management, visibility, and control. In many deployments, L3VPN services are offered over SD-WAN overlays to provide site-to-site connectivity with traffic segmentation, security, and performance guarantees. These L3VPN services leverage SD-WAN Secure Links, i.e. encrypted data plane tunnels established between SD-WAN edge nodes using mechanisms such as IPsec, to carry user traffic between endpoints.

This document describes the BGP mechanisms used to support such L3VPN deployments by enabling SD-WAN edge nodes to advertise underlay attributes, tunnel characteristics, and security association related attributes. These mechanisms enable dynamic tunnel selection, service-level steering, and secure endpoint discovery.

The SD-WAN usage model, including its deployment scenarios and BGP requirements, is detailed in [SD-WAN-BGP-USAGE] and not repeated here. This document focuses solely on the signaling extensions and encapsulation mechanisms required to support those scenarios in BGP.

1.2. SD-WAN Secure Links

[RFC9012] defines a BGP mechanisms that links routes (prefix and Next Hop) to a specific tunnels using a specific encapsulation. The SD-WAN Secure Links Topology uses a single hybrid logical link on a SD-WAN Peer to represent multiple underlay topology links. The SD-WAN peer distributes IPsec security association (IPsec-SA) [RFC4301] related information regarding the hybrid link or individual underlay links.

The traffic is routed via normal IPv4/IPv6 forwarding without any VPN addition. The SD-WAN Secure Links provides some link security for some simple cases of the three scenarios from [SD-WAN-BGP-USAGE] that do not require L3VPN addresses (Route Distinguisher (RD), prefix).

1.3. Conventions used in this document

The following acronyms and terms are used in this document:

Authorized BGP peer: A BGP peer that a given BGP speaker is permitted to exchange SD-WAN-related information with, based on local configuration or policy. This term is used in the context of SD-WAN edge discovery to indicate that not all discovered peers are automatically eligible for overlay connectivity.

C-PE (Customer Premises Equipment): A specific type of SD-WAN Edge deployed at the customer's edge. In this document, the terms C-PE and SD-WAN Edge are used interchangeably when referring to SD-WAN nodes that handle client route advertisement and secure tunnel establishment.

Cloud DC: Off-premises data centers that host applications and workloads, typically across multiple tenants or organizations, as defined in [Net2Cloud].

Controller: Refers to the SD-WAN Controller as defined in [SD-WAN-BGP-USAGE]. It is responsible for managing SD-WAN control-plane operations, including overlay path setup, teardown, policy enforcement, and route advertisement via BGP.

C-PE: Customer Premises Equipment that participates in the SD-WAN overlay network. As defined in [SD-WAN-BGP-USAGE], the term is used in this document to emphasize the role of a Provider Edge (PE) device functioning as the SD-WAN edge node, responsible for forming secure overlays across diverse underlay networks within the provider's infrastructure.

CPE-Based VPN: Virtual Private Secure network formed among C-PEs.

This is to differentiate such VPNs from most commonly used PE-based VPNs discussed in [RFC4364].

CPN: Customer Premises Network

IPsec-SA: IPsec Security Association [RFC4301].

MP_REACH_NLRI: A BGP Path Attribute used to advertise reachability information for multiple address families, as defined in [RFC4760]. This document uses MP_REACH_NLRI for carrying SD-WAN-specific NLRI.

SD-WAN: Software-Defined Wide Area Network, as defined in [MEF 70.1] and [MEF 70.2], refers to a policy-driven overlay connectivity service that operates over one or more underlay networks.

SD-WAN Edge: A network element that participates in the SD-WAN overlay by originating or receiving client routes, establishing secure underlay tunnels, and exchanging BGP-based control information.

SD-WAN Hybrid tunnel: A single logical tunnel that combines several links of different encapsulation into a single tunnel. This logical tunnel MAY exist as part of a SD-WAN Secure L3VPN or simply be a SD-WAN secure link for a flat network.

Secure Transport Connection: A transport layer security mechanism (e.g., IPsec, TLS, or SSL) layered under the BGP session to provide confidentiality, integrity, and authenticity of routing updates over untrusted networks.

TEA: Abbreviation for Tunnel Encapsulation Path Attribute [RFC9012]. Used in this document for brevity.

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 [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Framework of BGP SD-WAN Edge Discovery

This section provides a framework to describe the overall solution parts based on the reference diagram shown in Figure 1. The framework covers the following: client routes supported, example topologies, objectives of the SD-WAN Edge discovery, comparison of SD-WAN Secure VPNs with Pure IPsec VPNS, and what SD-WAN segmentation means in this SD-WAN context.

2.1. Supported Client Routes

Tunnel Encapsulation Attribute [RFC9012] MAY be attached to the routes carried in BGP UPDATES for IPv4 and IPv6 unicast routes (AFI/SAFI 1/1 and 2/1), as well as for IPv4 and IPv6 L3VPNs (AFI/SAFI 1/128 and 2/128), as specified in [RFC4364] and [RFC4659]. In the context of this specification, these NLRIs are used to signal SD-WAN tunnel information and endpoint properties associated with specific routes.

As described in Section 1.1, this document builds on the SD-WAN Secure VPN model defined in [SD-WAN-BGP-USAGE], [MEF-70.1], and [MEF70.2] to support L3VPN services. In this context, SD-WAN Secure L3VPNs leverage IPv4 and IPv6 L3VPNs ([RFC4364], [RFC4659]) extended with SD-WAN-specific mechanisms, including SD-WAN Hybrid Tunnel support and signaling of IPsec Security Association (IPsec-SA) attributes for underlay tunnel establishment.

As described in Section 1.2, the SD-WAN Secure Links model uses SD-WAN tunnels to create secure hybrid overlay connections between SD-WAN edge nodes. Unlike SD-WAN Secure L3VPNs, this application does not identify VPNs via the L3VPN NLRI. Instead, SD-WAN edge nodes use BGP to advertise unicast IPv4 and IPv6 routes (AFI/SAFI 1/1 and 2/1) along with the Tunnel Encapsulation Attribute (TEA) [RFC9012] containing an SD-WAN Hybrid Tunnel TLV (see Section 4). The SD-WAN Secure Links model also includes IPsec-SA metadata for underlay tunnels carried in BGP.

This document specifies the BGP mechanism for SD-WAN Secure L3VPN and SD-WAN Secure Links.

2.2. Topologies of SD-WAN Edge Discovery

This section describes how the topologies for SD-WAN Secure L3VPN and SD-WAN Secure Links can be served by the BGP SD-WAN logical Tunnel links.

2.2.1. SD-WAN Secure L3VPN Example Topology

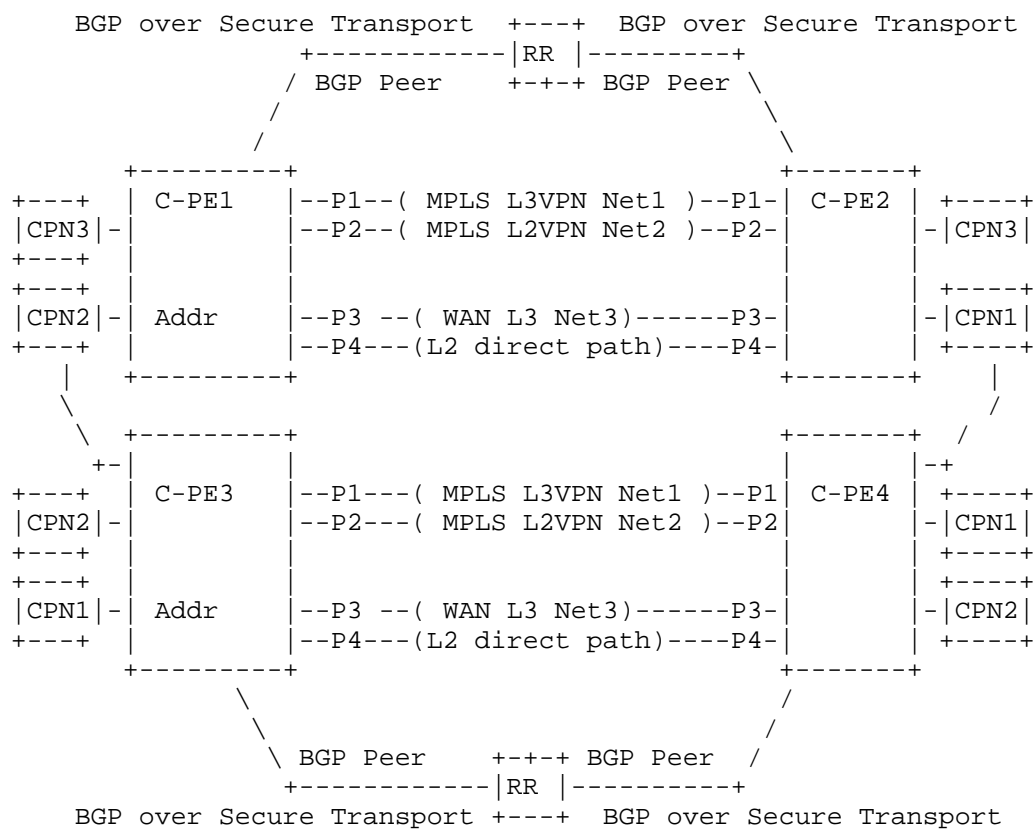
This section summarizes how a BGP-based control plane, using the BGP Tunnel Encapsulation Attribute [RFC9012], can support three deployment scenarios described in [SD-WAN-BGP-USAGE]. These scenarios focus on different ways encrypted tunnels are used in SD-WAN service delivery. While Section 1.1 lists five BGP-specific control-plane requirements, the following list highlights service-level SD-WAN deployment models that BGP can support through tunnel advertisement and associated metadata.

- * Homogeneous Encrypted SD-WAN,
- * Differential Encrypted SD-WAN, and
- * Private VPN PE based SD-WAN.

Based on [SD-WAN-BGP-USAGE], it is easy to see how Figure 1 aligns with Scenario 1 (Homogeneous Encrypted SD-WAN) and Scenario 2 (Differential Encrypted SD-WAN). Recasting Figure 1 as a logical BGP peering topology results in a BGP topology between PEs and C-PEs (customer premises equipment), as shown in Figure 2. Scenario 3 (Private VPNs based on SD-WAN) from [SD-WAN-BGP-USAGE] also corresponds to Figure 2.

Hybrid SD-WAN tunnel infrastructure requires that the PEs (C-PE1, C-PE2, C-PE3, and C-PE4) have an existing topology that the Hybrid SD-WAN tunnel overlays. These PEs use local policy to sending the appropriate traffic across the appropriate network based on local policy.

Hybrid SD-WAN Secure L3VPN: Sample Topology




Please note that RR MAY be one RR, but for clarity of diagram the RR has been displayed as two parts.

Figure 1: Hybrid SD-WAN Secure L3VPN

Hybrid SD-WAN BGP Mesh : Logical BGP Topology for SD-WAN

Logical topology for BGP peers



Note:

- SD-WAN Hybrid: SD-WAN Hybrid Tunnel

- For diagram clarity, the RR is shown as two separate nodes to illustrate its connectivity to all C-PEs. In practice, there may be only one RR. All BGP sessions to the RR are established over secure transport connections.

- The same CN depicted as connected to multiple C-PEs is to illustrate a multihoming scenario, where a single customer network maintains multiple connections to the SD-WAN fabric via different C-PEs.

Figure 2: Hybrid SD-WAN BGP Mesh

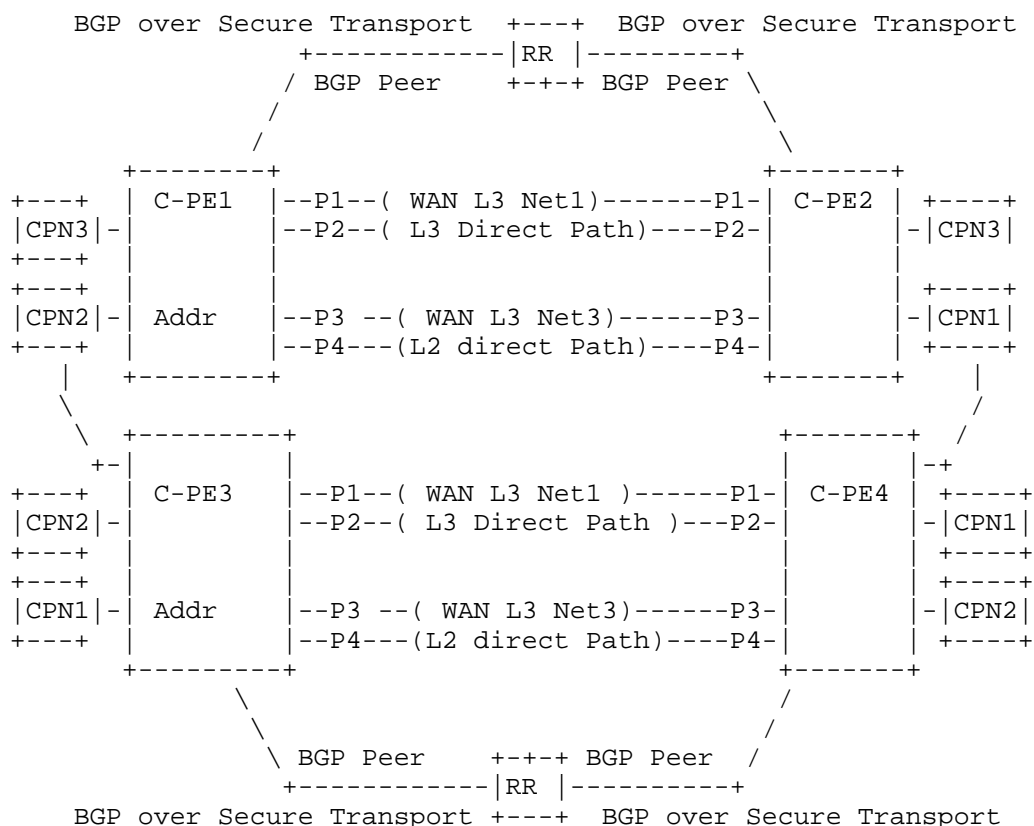
Note that the Figure 1 SD-WAN node BGP peers (C-PE) are connected in the underlay by both trusted VPNs and untrusted public networks. For trusted VPNs, IPsec Security associations MAY not be set-up. In some circumstances, some SD-WAN node peers MAY be connected only by untrusted public networks. For the traffic over untrusted networks, IPsec-SA must be established and maintained. If an edge node has network ports behind a NAT, the NAT properties need to be discovered by the authorized SD-WAN peers.

2.2.2. SD-WAN Secure Links Example Topology

Suppose the SD-WAN network topology from Figure 1 removes the L3VPN links. Instead the links are simply a combination of L3 WAN Networks (unsecured) and physically secure direct L2 and L3 links. The topology in Figure 1 becomes the underlay physical topology in Figure 3. The unsecured links need IPSec encryptions so the IPSec links must be configured. An SD-WAN Hybrid tunnel allows connections between the C-PEs (C-PE1, C-PE2, C-PE3, and C-PE) to be a single logical link.

Therefore, the logical topology in Figure 3, reduces to the SD-WAN logical topology shown in Figure 2.

Hybrid SD-WAN Secure Links: Sample Topology



Please note that RR MAY be one RR, but for clarity of diagram the RR has been displayed as two parts.

Figure 3: Hybrid SD-WAN Secure Links

2.2.3. RR connections to BGP Peers in SD-WAN

For both the SD-WAN Secure L3VPN network and the SD-WAN Secure Links, the BGP Peers are assumed to be connected to a RR via Secure Transport Connection. For an SD-WAN deployment with multiple RRs, it is assumed that there are Secure Transport Connection among those RRs.

How Secure Transport Connections are established between the BGP Peer and the RR, or between the multiple RRs is out of the scope of this document. The existing BGP UPDATE propagation mechanisms control the edge properties propagation among the RRs.

In some environments where the control plane communication to the RR is already secured using mechanisms such as those discussed in [RFC7018], IKE-less IPsec can be deployed to simplify data-plane IPsec-SA establishment among edge nodes..

2.3. The Objectives of SD-WAN Edge Discovery

The objectives of SD-WAN edge discovery are for SD-WAN edge nodes to discover the associated properties needed to establish secure overlay tunnels with their authorized BGP peers, where "authorized" refers to peers permitted by the SD-WAN domain policy as dictated by the controller [Net2Cloud]. The attributes to be propagated for the SD-WAN Secure L3VPN case are:

- * the SD-WAN (client) VPN information,
- * the attached client routes per VPN, and
- * the properties of the underlay networks over which the client routes can be carried.

Like any VPN network, client routes within specific SD-WAN VPNs can only be exchanged between authenticated SD-WAN peer nodes. Authentication of BGP sessions can be achieved using mechanisms such as TCP-AO [RFC5925], while additional protections (e.g., encryption and peer validation) may be provided by secure transport mechanisms like IPsec or TLS, depending on deployment practices.

The attributes to be propagated for the SD-WAN Secure L3 Links are:

- * the attached client routes and
- * the properties of the underlay networks over which the client routes can be carried.

The properties of the underlay network MAY include the following: IPsec-SA information, information needed for NAT transversal with IPsec, and link speeds.

Some SD-WAN peers are connected by both trusted VPNs and untrusted public networks. Some SD-WAN peers are connected only by untrusted public networks. For the traffic over untrusted networks, IPsec-SA must be established and maintained. For the trusted VPNs, IPsec Security associations MAY not be set-up. If an edge node has network ports behind a NAT, the NAT properties need to be discovered by the authorized SD-WAN peers.

2.4. Examples of SD-WAN Edge Discovery

Suppose the environment is Figure 1 with the logical SD-WAN Hybrid link topology of Figure 2. These topologies are set-up via configuration with IPsec-SA pre-configured. Suppose that C-PE1 and C-PE2 have 10 pre-shared keys to use on IPsec links. Currently P3 is using IPsec-SA ID-1. C-PE1 wants to receive traffic flowing from C-PE2 over the Hybrid SD-WAN links.

Step 1: Send Client Route with TEA to RR: Referring to the Figure 2, the C-PE1 routers send customer routes (L3VPN v4 route) to the RR with

- * NextHop set to self (C-PE1), and
- * attaching a Tunnel Encapsulation attribute specifying a Hybrid SD-WAN tunnel to the RR over a Secure Transport Connection

Step 2: RR forwards the Client route to C-PEs: Based on RR policy, the RR forwards routes to other BGP peers that support SD-WAN discovery for that AFI/SAFI (L3VPN v4). Using the Figure 2 example with the L3VPN for v4, RR would forward client routes with TEA of SD-WAN tunnel TLV and Next-Hop of C-PE1.

Step 3: Remote C-PE checks link: Prior to C-PE2 installing the L3VPN Unicast Routes (1/128), the C-PE2 MUST verify at least one of the underlay connections exist between C-PE1 and C-PE2. Local policy will define which link (or links) the traffic for this route goes over between C-PE1 and C-PE2.

Suppose for some reason the L2 link between C-PE1 and C-PE2 has encounter attacks, and the IPsec encryption must now run on links P3 and P4. C-PE1 detects the problem and to change the encryption on P3 and add encryption on P4. C-PE1 and C-PE2 MUST agree upon the next encryption on the list, and will send the IPsec-SA information in-band via BGP.

Step 1: C-PE1 Send Underlay Route with TEA to RR: C-PE1 sends two SD-WAN NLRIs in a BGP UPDATE with a TEA with SD-WAN Hybrid Tunnel TLV with an IPsec-SA-ID TLV with 4 IPsec-SA Identifiers (4, 5, 6, 7). The first SD-WAN NLRI contains Port P3's information (Port-local-ID P3, SD-WAN Color Gold (1), and SD-WAN Node-ID C-PE1). The second SD-WAN NLRI contains Port P4's information (Port-local-id P4, SD-WAN Color Gold (1), and SD-WAN Node-ID C-PE1).

Step 2: RR forwards Underlay route to C-PEs: The RR forwards the 2 underlay routes to C-PEs (C-PE1, C-PE2, C-PE3, C-PE4).

Step 3: C-PE2 receives the route and processes IPsec-SA. After receiving C-PE1 update, C-PE2 processes the update, and switches the IPsec-SA on Port 3 to IPsec-SA 4, and the IPsec-SA on P4 to IPsec-SA 4.

This simple example shows the value of rotating the pre-shared keys. Future IPsec-SA can also be set-up, negotiated, or rekeyed in the same manner.

The following question MAY occur to the observant reader:

- * Why is IPsec related information passed on a different AFI/SAFI?
- * Why do you do to support IPsec combined with NATs?
- * Is this the best way to support NATs?

The BGP SD-WAN Nodes can attach the TEA with a Hybrid SD-WAN TLV attached to the client routes (AFI/SAFI 1/1, 2/1, 1/128, 2/128) or attached to the SD-WAN NLRI (AFI/SAFI 1/74 or 2/74). The SD-WAN NLRI allows the passing of IPsec information on a unique AFI/SAFI. How implementation prioritizes the processing of client routes versus underlay routes (SD-WAN NLRI) is an implementation matter, and out of scope for this document.

2.5. Comparing SD-WAN Secure L3VPN with Pure IPsec VPN

This section highlights the benefits of distributing IPsec-SA attributes via a new BGP AFI/SAFI using the Tunnel Encapsulation Attribute (TEA) [RFC9012], rather than relying solely on traditional IPsec signaling mechanisms such as IKEv2 [RFC7296]. In BGP-controlled SD-WAN environments where BGP already operates over a secure control plane (e.g., using TLS or TCP-AO), reusing BGP as the signaling channel simplifies deployment by avoiding the mutual authentication and multi-step negotiation process required by IKEv2, thereby reducing complexity and operational overhead.

Establishing an IPsec-SA between two untrusted nodes typically requires the following configurations and message exchanges:

IPsec IKEV2: Messages are sent to authenticate with each other.

Establish IPsec-SA: Establishing the IPsec-SA requires the following set-up

- Local key configuration,
- Setting the Remote Peer address,

- Information from IKEv2 Proposal directly sent to peer (encryption method, integrity sha512, etc.), and
- Transform set.

Discovery of attached client routes is achieved by:

- Running routing protocol within each IPsec-SA.
- If multiple IPsec-SAs between two peer nodes are established to achieve load sharing, each IPsec tunnel needs to run its own routing protocol to exchange client routes attached to the edges.

Set-up Access List or Traffic Selector: Access Lists permit specific Local-IP1, Remote-IP2, and other IPsec parameters.

In a BGP-controlled SD-WAN network over hybrid MPLS VPN and public internet underlay networks, all edge nodes and RRs are already connected by private secure paths. The RRs have the policies to manage the authentication of all peer nodes. More importantly, when an edge node needs to establish multiple IPsec tunnels to many edge nodes, all the management information can be multiplexed into the secure management tunnel between RR and the edge node operating as a BGP peer. Therefore, the amount of authentication in a BGP-Controlled SD-WAN network can be significantly reduced.

Client VPNs are configured via VRFs, just like the configuration of the existing MPLS VPN. The IPsec equivalent traffic selectors for local and remote routes are achieved by importing/exporting VPN Route Targets. The binding of client routes to IPsec-SA is dictated by policies. As a result, the IPsec configuration for a BGP controlled SD-WAN (with mixed MPLS VPN) can be simplified in the following manner:

- * The SD-WAN controller has the authority to authenticate edges and peers so the Remote Peer association is controlled by the SD-WAN Controller (RR).
- * The IKEv2 proposals (including the IPsec Transform set) can be sent directly to peers, or incorporated in a BGP UPDATE.
- * The BGP UPDATE announces the client route reachability through the SDWAN hybrid tunnels. A SDWAN hybrid tunnel combines several other tunnels into a single logical tunnel. The SD-WAN Hybrid tunnel implementations insure that all tunnels within are either running over secure network links or secured by IPsec.

- * Importing/exporting Route Targets under each client VPN (VRF) achieves the traffic selection (or permission) among clients' routes attached to multiple edge nodes.

Note: The web of SDWAN hybrid tunnels in a network is denoted in this document as an SD-WAN underlay. BGP passes information about the SDWAN hybrid tunnels between BGP peers by passing an SD-WAN Underlay NLRI paired with the tunnel encapsulation attribute (TEA) with an SDWAN tunnel type SD-WAN-Hybrid TLV.

Also, note that with this method there is no need to run multiple routing protocols in each IPsec tunnel.

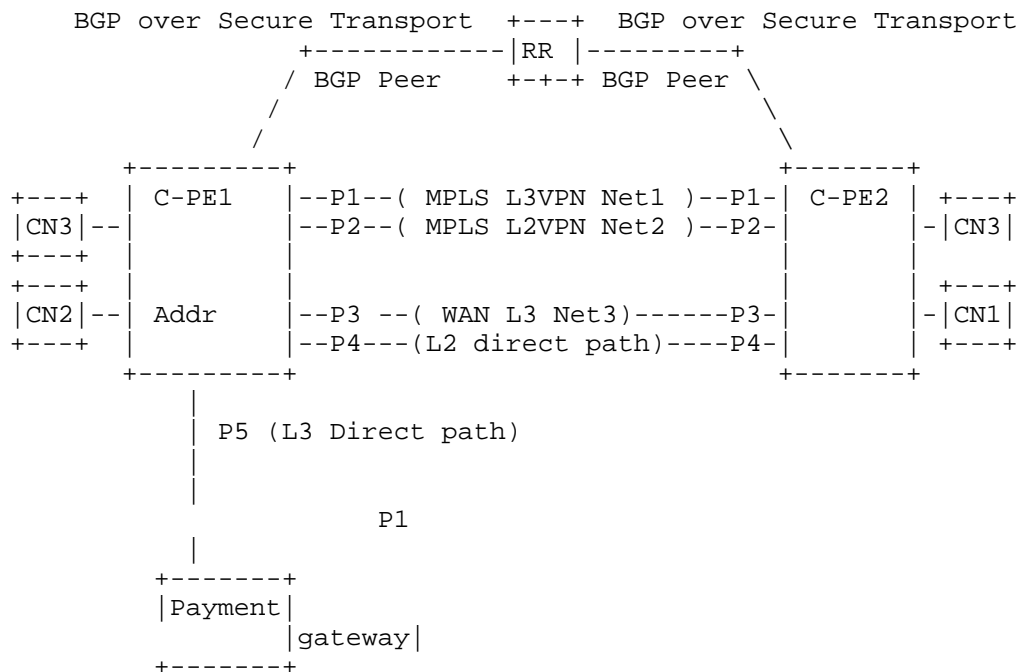
2.6. SD-WAN Segmentation, Virtual Topology and Client VPN

In SD-WAN Secure L3VPN deployments, SD-WAN Segmentation is a frequently used term which refers to partitioning a network into multiple subnetworks, just like MPLS VPNs. SD-WAN Segmentation is achieved by creating SD-WAN virtual topologies and SD-WAN VPNs.

An SD-WAN virtual topology consists of a set of edge nodes and the tunnels (a.k.a. underlay paths) interconnecting those edge nodes. These tunnels forming the underlay paths can be IPsec tunnels, or MPLS VPN tunnels, or other tunnels. Figure 4 (top) illustrates an SD-WAN L3VPN underlay topology and Figure 4 (bottom) show the same topology as the virtual topology based on SD-WAN Links.

An SD-WAN VPN is configured in the same way as the VRFs of an MPLS VPN. One SD-WAN client VPN can be mapped to multiple SD-WAN virtual topologies. A Route Target is used to differentiate between the SD-WAN VPNs. For example, in figure 4 below, the Payment-Flow on C-PE2 is only mapped to the virtual topology of C-PEs to/from Payment Gateway, whereas other flows can be mapped to a multipoint-to-multipoint virtual topology.

The SD-WAN Controller governs the policies for mapping client VPNs to SD-WAN virtual topologies. Each SD-WAN edge node may support multiple VPNs, and Route Targets are used to differentiate among them. For example, on C-PE2, the "Payment-Flow" VPN may be mapped only to a specific virtual topology involving C-PEs that connect to the Payment Gateway, while other VPNs may be associated with a broader multipoint-to-multipoint topology.



Tunnel exist C-PE2 port P4 on L2 direct path with encryption to P1 port on Payment gateway.




Figure 4: SD-WAN Virtual Topology and VPN

3. BGP SD-WAN Mechanisms

The BGP mechanisms have two functions:

Advertise Client routes with Hybrid SD-WAN tunnel: A BGP speaker supporting SD-WAN re-advertises routes received from client routers with the Next_HOP attributes set to its own IP address, as explicitly configured [RFC4271], and include BGP attribute indicating the SD-WAN Hybrid Tunnel. Clients routes MAY be advertised using the following AFI/SAFIs: Unicast IPv4/IPv6(1/1, 2/1) and L3VPN IPv4/IPv6 (1/128, 2/128). The term "next hop self" means the routers sets the next Hop Address to an address indicating the BGP Peer. The SD-WAN tunnel indication can be conveyed using either the Encapsulation Extended Community (Encap-EC) or the Tunnel Encapsulation Attribute (TEA).

Advertise Underlay Routes (SD-WAN NLRIs) with Hybrid SD-WAN TEA: A BGP speaker advertises SD-WAN NLRI for IPv4/IPv6 (AFI/SAFI 1/74 or 2/74) with the NEXT_HOP attribute set to the local address of the advertising speaker, as explicitly configured [RFC4271], and includes a BGP attribute indicating the Hybrid Tunnel. The SD-WAN NLRI identifies the port (or ports) within the Hybrid SD-WAN Tunnel for which the BGP speaker is advertising encapsulation or IPsec-SA related information via the Hybrid SD-WAN TEA. The Hybrid SD-WAN TEA contains IPsec-SA and, optionally, NAT-related information.

This section describes the Hybrid SD-WAN tunnel, the SD-WAN NLRIs, the new sub-TLVs for SD-WAN Tunnel IPsec-SA, sub-TLVs for Port attributes, the procedures for the client routes, the procedures for underlay routes, error handling, and considerations for managing SD-WAN technologies.

3.1. SD-WAN Hybrid Tunnel Encoding

Name: SD-WAN Hybrid Tunnel


Code: 25 (IANA assigned)

Description: The SD-WAN Hybrid Tunnel identifies a virtual tunnel that overlays a path across a set of underlay links between two BGP peers. These underlay links may use various technologies (e.g., MPLS, Layer 2 direct connections, or Layer 3 public Internet). The term hybrid reflects that different types of underlay links can be used simultaneously.

Encoding: Per [RFC9012], the following two BGP attributes that MAY

encode a Tunnel Encapsulation attribute information: the Tunnel Encapsulation Attribute, and the Encapsulation Extended Community (Encap-EC) as a "barebones" tunnel identification. The encoding for the SD-WAN Hybrid tunnel is described for both BGP attributes.

SD-WAN Encoding in Encap-EC



Encapsulation Extended Community

Figure 5: SD-WAN Hybrid tunnel encoding in Encap-EC

The NextHop Field in the BGP update is the tunnel egress Endpoint, and this SHOULD be set to the BGP Peer Address for the SD-WAN Peer.

SD-WAN Encoding in TEA




Figure 6: SD-WAN Hybrid Value Field

No Sub-TLVs for the Encap-EC of the Hybrid SD-WAN Tunnel

[RFC9012]: When Encap-EC with Tunnel Type = 25 is attached to a client route, the detailed SD-WAN tunnel attributes, particularly those related to IPsec parameters and keying material, are not included in the same BGP UPDATE message. Instead, they are advertised separately using the SD-WAN NLRI, as described in Section 4.2 and 4.3. The SD-WAN NLRI is originated using the loopback address of the C-PE, rather than the client route. The remote BGP speaker uses this loopback address to associate the client route with the corresponding SD-WAN Hybrid Tunnel. This

separation allows for independent advertisement rates and avoids bloating BGP UPDATE messages with the large amount of data required for IPsec-SA, cryptographic keys, and related parameters.

Sub-TLVs for Hybrid SD-WAN Tunnel (Type 25) in the Tunnel Encapsulation Attribute (TEA): When the Hybrid SD-WAN Tunnel (Type 25) is used within the Tunnel Encapsulation Attribute for client routes, all sub-TLVs defined in [RFC9012] are valid, along with the additional sub-TLVs specified in Section 4.2 and 4.3 of this document. In this case, detailed underlay tunnel attributes, such as IPsec-related parameters, are included directly in the same BGP UPDATE as the client route. As a result, there is no need for a separate UPDATE message associated with the C-PE loopback address. However, this approach means that any changes to underlay attributes (e.g., IPsec keys or cryptographic parameters) necessitate re-advertising the client route with an updated TEA, which can increase both the frequency and size of BGP UPDATE messages.

Validation Procedure: The validation procedure for the SD-WAN tunnel TLV has the following components:

- 1) validation of tunnel TLV encoding [RFC9012],
- 2) Check that sub-TLVs are valid for NLRI (see Table 1), and
- 3) Egress Tunnel End Point Check: validate that the tunnel egress endpoint (as carried in the Tunnel Egress Endpoint sub-TLV [RFC9012]) is a reachable IP address based on the BGP next-hop resolution rules.

Prior to installing a route with a SD-WAN tunnel as an active route, the BGP peer installing the route MUST also validate that the SD-WAN tunnel and underlay links are active.

Table 1

Client Routes AFI/SAFI = 1/1, 2/1, 1/128, 2/128

Underlay Routes AFI/SAFI = 1/74 and 2/74

sub-TLV	Code	Client Routes	Underlay Routes
-----	----	-----	-----
Encapsulation	1	not valid	not valid
Protocol	2	not valid	not valid
Color	3	valid *1	not valid *2
Load-Balancing Bk	5	not valid	not valid *2
Tunnel Egress EP	6	Info rqd *5	required
DS Field	7	not valid	not valid *2
UDP Dest. Port	8	not valid	not valid *2
Embedded Label H.	9	not valid	not valid *2
MPLS label Stack	10	not valid	not valid *2
Prefix-SID	11	not valid	not valid *2
Preference	12	not valid	not valid *2
Binding SID	13	not valid	not valid *2
ENLP	14	not valid	not valid *2
Priority	15	not valid	not valid *2
SPI/SI	16	not valid	not valid *2
SRv6 Binding SID	20	not valid	not valid *2
IPsec-SA-ID	64	valid *3	valid *3
Extended Port Attr	65	not valid	valid *4
Underlay Type	66	not valid	valid *4
IPsec-SA Rekey Cnt	67	valid *3	valid *3
IPsec Public Key	68	valid *3	valid *3
IPsec-SA Proposal	69	valid *3	valid *3
Simplified IPsec-SA	70	valid *3	valid *3

Figure 7: Table 1: sub-TLV list

Notes

- * *1 - For client traffic, the Color sub-TLV defined in this document must be validated using the same procedures specified in [RFC9012] for the Color Extended Community. When both the Color Extended Community and the Color sub-TLV are present, the value in the Color Extended Community [RFC9012] takes precedence and must be used for forwarding and policy decisions.

- * *2 - The listed Sub-TLVs are not valid when used with Underlay route advertisements. Future extensions may define their use in that context, but such extensions are outside the scope of this document.
- * *3 - See Sections 4.3 (encoding), 4.4 (client route validation), and 4.5 (underlay route validation) provides content processing and validation procedures for underlay routes
- * *4 - See Sections 4.3 (encoding), and 4.6 (error handling for malformed sub-TLVs or incorrect NLRI association).
- * *5 - Per [RFC9012] Section 4.1, when TEA is attached to a client route UPDATE, the Tunnel Egress Endpoint is derived from the BGP NextHop attribute.

3.2. SD-WAN Underlay UPDATE

The Edge BGP Peer using BGP SD-WAN discovery sends the hybrid SD-WAN NLRI with the SD-WAN Hybrid tunnel attribute to advertise the detailed properties associated with the public facing WAN ports and IPsec tunnels. The Edge BGP Peer sends this information to its designated RR via the Secure Transport Connection. Each BGP UPDATE message with a SD-WAN Underlay NLRI MUST contain a Tunnel Encapsulation Attribute (TEA) for a Hybrid Tunnel type. The TEA can include sub-TLVs for Extended Port attribute (see section 3.3.6) or IP Sec information (see section 8). The IPsec information sub-TLVs include: IPsec-SA-ID, IPsec-SA Nonce, IPsec Public Key, IPsec-SA Proposal, and Simplified IPsec-SA.

3.2.1. The NLRI for SD-WAN Underlay Tunnel Update

A new NLRI SAFI (SD-WAN SAFI=74) is introduced within the MP_REACH_NLRI Path Attribute of [RFC4760] for advertising the detailed properties of the SD-WAN tunnels terminated at the edge node:

```

+-----+
|   Route Type   | 2 octets
+-----+
|   Length       | 2 octets
+-----+
| Type Specific  |
~ Value (Variable) ~
|
+-----+
```

Figure 8: SD-WAN NLRI Encoding

where:

Route (NLRI) Type: 2 octet value to define the encoding of the rest of the SD-WAN the NLRI.

Length: 2 octets indicating the length of the Value field, expressed in bits, following the NLRI encoding convention defined in [RFC4760], Section 3.

This document defines the following SD-WAN Route type:

NLRI Route Type = 1 (SD-WAN Tunnel Endpoint NLRI): For advertising the detailed properties of the SD-WAN tunnels terminated at the edge, where the transport network port can be uniquely identified by a tuple of three values (Port-Local-ID, SD-WAN-Color, SD-WAN-Node-ID). The SD-WAN NLRI Route Type =1 has the following encoding:

+-----+	
Route Type = 1	2 octets
+-----+	
Length	2 octets
+-----+	
Port-Local-ID	4 octets
+-----+	
SD-WAN-Color	4 octets
+-----+	
SD-WAN-Node-ID	4 or 16 octets
+-----+	

Figure 9: SD-WAN NLRI Route Type 1

Length: The value of the Length field for Route-Type 1 MUST be either 12 octets (when the SD-WAN-Node-ID is an IPv4 address) or 24 octets (when the SD-WAN-Node-ID is an IPv6 address). Any other value is invalid, and the NLRI MUST be treated as malformed and discarded.

Port-local-ID: SD-WAN edge node Port identifier, which is locally significant. If the SD-WAN NLRI applies to multiple WAN ports, this field is zero.

SD-WAN-Color: identifies a group of SD-WAN tunnels that may span

multiple SD-WAN edges co-located at the same site. This value correlates with the Color Extended Community attached to client routes. The receiving BGP speaker selects SD-WAN tunnels whose SD-WAN-Color matches the Color Extended Community in the client route when determining which underlay tunnel(s) to use. If the SD-WAN-Color represents all tunnels at a site, it effectively serves as a site-level identifier. If no matching SD-WAN-Color is found, the client route may not be forwarded over any SD-WAN tunnels.

SD-WAN Node ID: This field carries the IPv4 or IPv6 address of the SD-WAN edge node (C-PE). For IPv4 SD-WAN NLRI (AFI/SAFI 1/74), this field contains a 4-octet IPv4 address representing a /32 host address. For IPv6 SD-WAN NLRI (AFI/SAFI 2/74), this field contains a 16-octet IPv6 address representing a /128 host address. The SD-WAN Node ID identifies the loopback address used by the SD-WAN edge node to advertise its tunnel attributes.

3.2.2. Validation of SD-WAN NLRI

Upon receiving an SD-WAN NLRI Route-Type 1, the following validation steps MUST be performed:

The Length field MUST contain a value of either 12 or 24 octets, as defined in Section 4.2.1. Any other value renders the NLRI malformed and it MUST be discarded.

If Length = 12, the SD-WAN Node-ID field MUST contain exactly 4 octets, representing an IPv4 address. If Length = 24, the SD-WAN Node-ID field MUST contain exactly 16 octets, representing an IPv6 address.

The SD-WAN Node-ID MUST be a valid unicast address.

Implementations MAY apply additional local policy checks (e.g., verifying whether the advertising BGP speaker is authorized to advertise SD-WAN NLRIs), but these are outside the scope of NLRI field validation itself.

If the local policy check fails, the NLRI SHOULD be discarded without affecting the BGP session.

3.2.3. BGP Path Attributes attached to SD-WAN NLRI

The Path Attributes attached to the SD-WAN NLRIs apply to the WAN-facing tunnel endpoints being advertised, not to client routes. These attributes describe properties of the WAN ports (e.g., encapsulation, transport role, or color) that may be used in establishing SD-WAN overlay tunnels between edge nodes. Client routes, which represent customer prefixes, are propagated using separate BGP NLRIs (e.g., IPv4/IPv6 unicast or L3VPN), with their own associated Path Attributes. The SD-WAN NLRI and client route NLRI are independent but may be correlated by the receiving BGP speaker for tunnel selection and service mapping.

3.3. IPsec-SA Property Sub-TLVs

The IPsec-SA Property Sub-TLVs defined in this section are used to signal IPsec-SA parameters for Hybrid SD-WAN tunnels as defined in this document. While these Sub-TLV formats could potentially be reused in other applications that require IPsec-SA signaling over BGP, this document defines their semantics and behavior specifically within the SD-WAN Edge Discovery framework.

If any sub-TLV is malformed, the implementation MUST follow the procedure in [RFC9012] in section 13.

To support key rotation (e.g., updating IPsec keys or parameters), the SD-WAN NLRI (identified by Port-Local-ID, SD-WAN-Color, and SD-WAN-Node-ID) can be re-advertised via a BGP UPDATE message containing updated IPsec-SA information. This mechanism enables rapid distribution of new keys without requiring separate key negotiation protocols.

3.3.1. IPsec-SA-ID Sub-TLV

The IPsec-SA-ID Sub-TLV is used to reference one or more previously established IPsec-SAs between SD-WAN nodes. This Sub-TLV carries one or more 32-bit Security Parameter Index (SPI) values assigned at the receiving node (i.e., the inbound SPI). When combined with the SD-WAN Node-ID (which identifies the tunnel endpoint address), each SPI uniquely identifies an existing IPsec-SA, consistent with the SA identification described in [RFC4301], Section 4.1.

Multiple SPIs MAY be included within the Sub-TLV to reference multiple pre-established IPsec-SAs available for the SD-WAN overlay. This enables advertisement of SA updates, key rotations, or operational state changes without resending full SA parameter sets, thereby significantly reducing the size of BGP UPDATE messages and allowing pairwise IPsec rekeying to proceed independently for each SA.

Sub-TLV Name: IPsec-SA-ID

Sub-TLV Code: 64 (IANA assigned)

Sub-TLV Encoding:

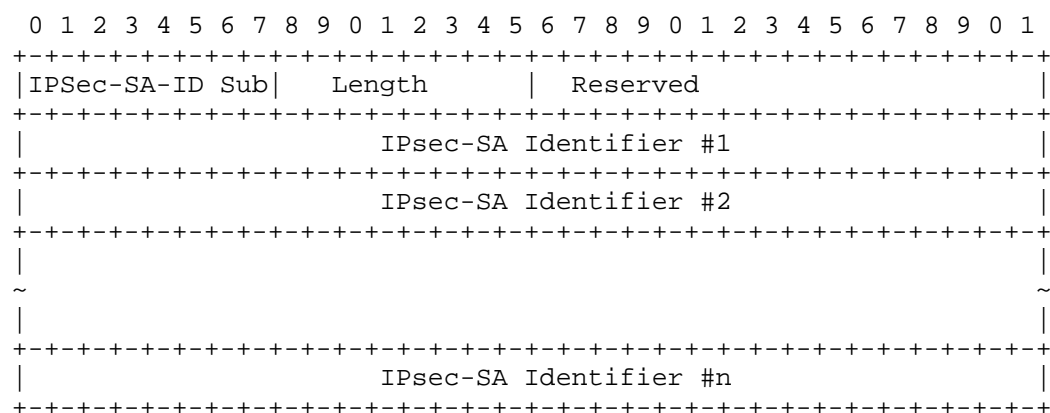


Figure 10: IPsec-SA-ID sub-TLV

where:

- * IPsec-SA-ID Sub-Type (8 bits): 64(IANA Assigned).
- * length (8 bits): Specifies the total length in octets of the value field (not including the Type and Length fields). For the IPsec-SA-ID Sub-Type, the length SHOULD be the 2 + 4 *(number of IPsec-SA IDs) .
- * Reserved: Reserved for future use. In this version of the document, the Reserved field MUST be set to zero and MUST be ignored upon receipt. Received values MUST be propagated without change.
- * Value field: The value field consists of a sequence of IPsec-SA SPIs, each 4 octets long. As shown in the figure above, n IPsec-SAs are attached in the IPsec-SA-ID sub-TLV:

- IPsec SA Identifier #1: A 4 octet SPI for a pre-established IP security association.
- IPsec SA Identifier #2: A 4 octet SPI for a pre-established IP security association.
- IPsec SA Identifier #n: A 4 octet SPI for a pre-established IP security association.

Sub-TLV Error Handling: The Length field of the IPsec-SA-ID Sub-TLV MUST be a non-zero multiple of 4 octets. Any other value is considered a malformed Sub-TLV. Error handling for malformed Sub-TLVs follows [RFC9012]

3.3.2. IPsec-SA Rekey Counter Sub-TLV

The IPsec-SA Rekey Counter Sub-TLV provides the rekey counter for a security association (identified by IPsec-SA Identifier).

Sub-TLV Name: IPsec-SA ReKey Counter - Rekey Counter for a IPsec-SA

Sub-TLV Code: 67 (IANA assigned)

Sub-TLV Encoding:

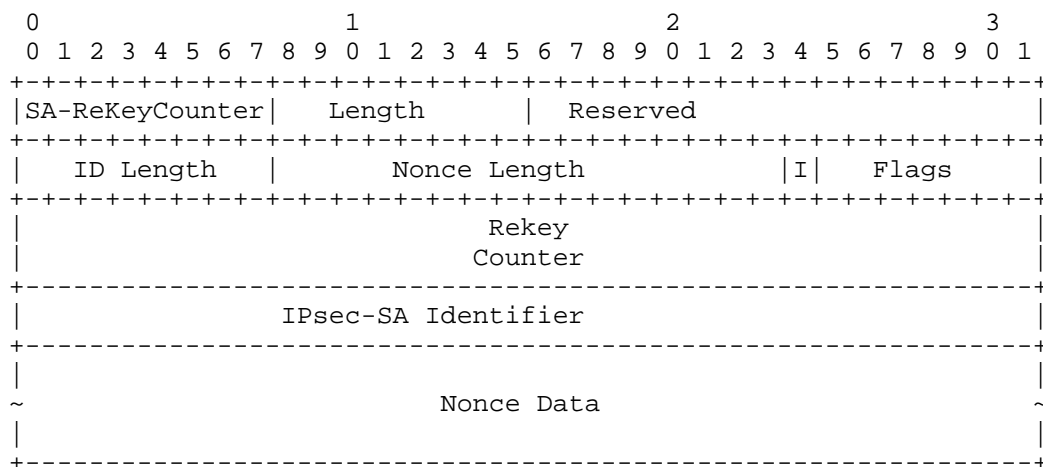


Figure 11: IPsec-SA Rekey Counter Sub-TLV Value Field

where:

- * SA-ReKeyCounter (IPsec-SA Rekey Counter) Sub-Type (8 bits): 67 (IANA assigned)

- * length (8 bits): Specifies the total length in octets of the value field (not including the Type and Length fields). The total length is variable with the value equal to 18 plus Nonce length.
- * Reserved: Reserved for future use. In this version of the document, the Reserved field MUST be set to zero and MUST be ignored upon receipt. Received values MUST be propagated without change.
- * ID Length (8 bits): indicates the length in octets of SA-Identifer (SA-SPI). This length SHOULD be 4 octets.
- * Nonce Length (16 bits): indicates the length, in octets, of the Nonce Data. The value MUST be a non-zero multiple of 4 (i.e., the Nonce Data length MUST be a multiple of 32 bits)[RFC7296].
- * I Flag: when set to 1, the I-flag indicates that the communication being established is new. when set to 0, it signals that the communication is a continuation of an existing session.
- * Flags (7 bits): Reserved for future use. In this version of the document, the Reserved field MUST be set to zero and MUST be ignored upon receipt. Received values MUST be propagated without change.
- * Rekey Counter (64 bits): the number of key updates or rekeys that have occurred. Each time a key is rotated or replaced, the ReKey Counter is incremented.
- * IPsec-SA Identifier (IPsec-SA-ID): identifies the SPI assigned to a specific IPsec-SA terminated at the SD-WAN edge node. The length of this field is specified in ID Length. For this specification, the length MUST be 4 octets. Other lengths are outside the scope of this document.
- * Nonce Data: a random or pseudo-random number for preventing replay attacks.

Sub-TLV Error Handling: The IPsec-SA Rekey Counter Sub-TLV is considered malformed under any of the following conditions:

The total Sub-TLV Length is less than the sum of ID Length, Nonce Length, and 4 octets for the Rekey Counter.

The ID Length field does not match the actual length of the ID field.

The Nonce Length field is zero or not a multiple of 4.

Malformed Sub-TLVs are handled according to [RFC9012]; they MUST be ignored and skipped during parsing.

3.3.3. IPsec Public Key Sub-TLV

The IPsec Public Key Sub-TLV provides the Public Key exchange information and the life span for the Diffie-Hellman Key. The encoding is shown in the figure below:

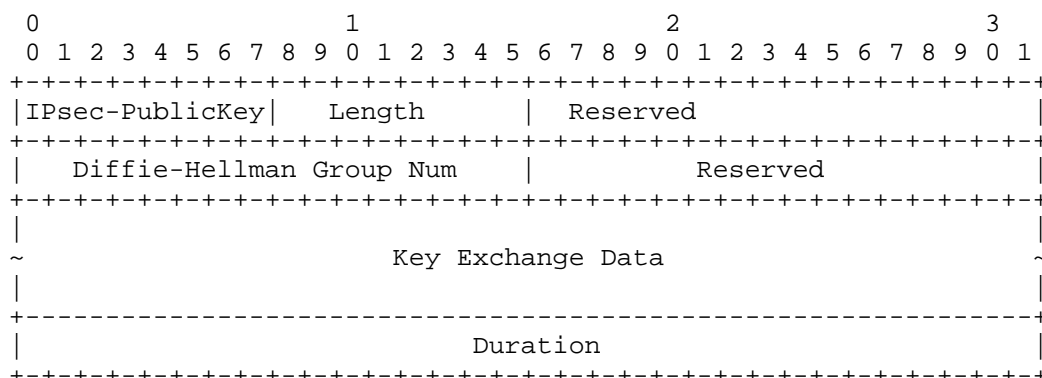


Figure 12: IPsec-SA Public Key Sub-TLV Value Field

where:

IPSec-PublicKey Sub-Type (8 bits): 68 (IANA assigned)

length (8 bits): Specifies the total length in octets of the value field (not including the Type and Length fields). The total length is variable with the length being 10 + the Key Exchange Data length.

Diffie-Hellman Group Num (16-bits): identifies the Diffie-Hellman group used to compute the Key Exchange Data. Details on Diffie-Hellman group numbers can be found in Appendix B of IKEv2 [RFC7296] and [RFC5114].

The Key Exchange data: This refers to a copy of the sender's Diffie-Hellman public value. The length of the Diffie-Hellman public value is defined for MODP groups in [RFC 7296] and for ECP groups in [RFC 5903].

Duration (32 bits): a 4-octet value specifying the life span of the Diffie-Hellman key in seconds.

An IPsec Public Key Sub-TLV is considered malformed if any of its fields do not conform to the encoding rules specified above. Malformed Sub-TLVs are handled according to [RFC9012] and MUST be ignored.

3.3.4. IPsec-SA Proposal Sub-TLV

The IPsec-SA Proposal Sub-TLV is used to advertise a set of cryptographic parameters that define the proposal for establishing an IPsec-SA. A proposal consists of one or more transform types, where each transform specifies a particular cryptographic function (such as encryption or integrity) and the corresponding algorithm to be used. This structure follows the same model as IKEv2 Proposals defined in [RFC7296] Section 3.3.

Sub-TLV Name: IPsec-SA Proposal - Indicates IPsec Transform Attributes

Sub-TLV Code: 69 (IANA assigned)

Each transform includes:

- A Transform Type, which identifies the function being specified (e.g., encryption, integrity).
- A Transform ID, which specifies the algorithm for that function.
- Optional Transform Attributes, which provide additional algorithm-specific parameters when necessary.

The encoding is shown below:>

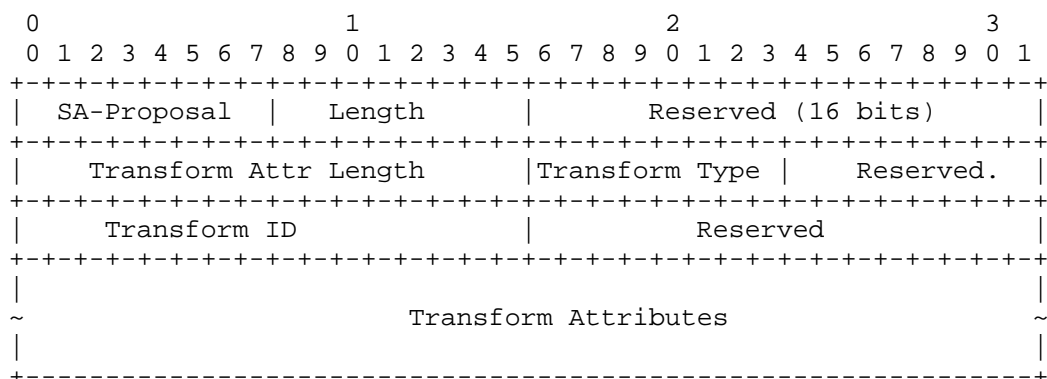


Figure 13: IPsec-SA Proposal Sub-TLV Value Field

where:

IPSec-SA-Proposal Sub-Type (8 bits): 69 (IANA assigned)

length (8 bits): Total length of the value field in octets (not including Type and Length fields). This equals 10 + the Transform attribute length.

Reserved (16 bits): reserved for future use. These bits are ignored upon receipt and set to zero when transmitted.

Transform Attr Length (16 bits): length of the Transform Attributes field in octets.

Transform Type (8 bits): The function being specified. Transform Type values are defined in [RFC7296] Section 3.3.2 and IANA IKEv2 Transform Type registry. Valid types include ENCR (1), PRF (2), INTEG (3), DH (4), and ESN (5).

Reserved (8 bits): Reserved for future use. MUST be set to zero when transmitted and ignored upon receipt.

Transform ID (16 bits): Identifies the algorithm for the corresponding Transform Type, as defined in [RFC7296] Section 3.3.2.

Transform Attributes: Optional algorithm-specific parameters, encoded as defined in [RFC7296] Section 3.3.5.

The Transform Attributes field may be omitted if no additional parameters are required for the selected algorithm.

Multiple IPsec-SA Proposal Sub-TLVs MAY be included to describe multiple transform types for the same SA proposal. Collectively, these Sub-TLVs define the full proposal for an IPsec-SA between SD-WAN edge nodes.

Sub-TLV Error Handling: An IPsec-SA Proposal Sub-TLV is considered malformed if:

- The Length field does not match the actual length of the Value field.
- The Transform Attribute Length field is inconsistent with the total Sub-TLV Length.
- Any field value falls outside its valid range as specified in [RFC7296].

Malformed Sub-TLVs MUST be handled according to [RFC9012] and ignored during parsing. Additional content checks for the IPsec-SA Proposal Sub-TLV are described in Section 4.4 (for client routes) and Section 4.5 (for underlay routes).

3.3.5. Simplified IPsec-SA Sub-TLV

The Simplified IPsec-SA Sub-TLV provides a compact way to signal pre-configured IPsec-SA parameters for deployments where full transform negotiation (e.g., via IKEv2) is not supported or not necessary. In such deployments, SD-WAN edge nodes are provisioned (e.g., via SD-WAN controller or management system) with a common set of agreed security profiles, including allowed transforms and algorithms. This Sub-TLV signals which profile entry is to be used for a given SA instance.

Sub-TLV Name: Simplified IPsec-SA

Sub-TLV Code: 70 (IANA assigned)

Sub-TLV Encoding:

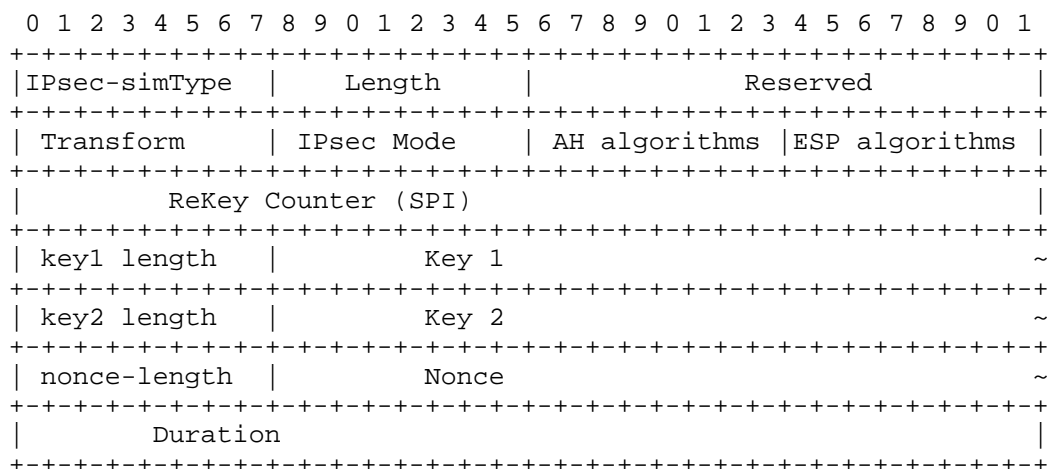


Figure 14: Simplified IPsec-SA Sub-TLV

where:

Length (8 bits): variable in octets (based on key length)

Reserved (16 bits): Reserved for future use. These bits SHOULD be set to zero on transmission and MUST be ignored on receipt.

Transform (8 bits):

- * Transform = 1 means AH,
- * Transform = 2 means ESP, or
- * Transform = 3 means AH+ESP.

All other transform values are invalid unless specified by future specifications.

IPsec Mode (8 bits):

- * Mode = 1 indicates that the Tunnel Mode is used.
- * Mode = 2 indicates that the Transport mode is used.

Only Mode values 1 and 2 are valid in this document. All other Mode values are considered invalid unless specified by future specifications.

AH algorithms (8 bits): Specifies the AH authentication algorithm to

be used. The values are defined in [RFC4305] and its updates (e.g., [RFC8221]). While an SD-WAN edge node may be capable of supporting multiple AH algorithms, this field carries only a single algorithm value for the specific SA instance. The selection of which algorithms are supported across peers is determined via SD-WAN controller provisioning or management policy. No in-band negotiation of multiple algorithms is performed using this field.

ESP algorithms (8 bits): Specifies the ESP encryption algorithm, as defined in [RFC4305], [RFC8221], and their updates. Like AH Algorithm, only a single algorithm value is carried per SA instance, with acceptable algorithms coordinated by provisioning or policy.

Rekey Counter (Security Parameter Index) (4 octet): indicates the count for rekeying.

key1 length (8 bits): indicates the IPsec public key 1 length

Public Key 1: IPsec public key 1

key2 length (8 bits): indicates the IPsec public key 2 length

Public Key 2: IPsec public key 2

nonce-length (8 bits): indicates the Nonce key length

Nonce: IPsec Nonce

Duration (32 bits): specifying the security association (SA) life span in seconds.

A Simplified IPsec-SA Sub-TLV is considered MALFORMED if any of its fields are not properly encoded, do not conform to the specified value ranges above, or contain invalid field lengths. Per [RFC9012], any MALFORMED Sub-TLV MUST be ignored, and processing continues with the remaining Sub-TLVs in the TEA.

3.3.6. Extended Port Attribute Sub-TLV

The Extended Port Attribute Sub-TLV advertises NAT-related properties associated with a public Internet-facing WAN port on an SD-WAN edge node. This information enables peer SD-WAN nodes to establish secure tunnels even when one or both peers are behind NAT devices. An SD-WAN edge node may query a STUN server (Session Traversal Utilities for NAT [RFC8489]) to determine its NAT properties, including its public IP address and public port number. These properties are then

advertised to peer nodes using the Extended Port Attribute Sub-TLV.

In SD-WAN deployments, NAT devices may exist at one or both ends of the tunnel path. The possible deployment scenarios include:

- * Only one SD-WAN edge node is located behind a NAT device, while its peer is directly reachable.
- * Both SD-WAN edge nodes are behind NAT devices (symmetric or independent NATs).
- * The external address and port assigned to an edge node may change dynamically, either due to ISP address allocation or when traversing NAT devices that use dynamic address pools.

Because an SD-WAN edge node may have multiple WAN ports with independent NAT characteristics, the NAT properties are associated with individual WAN ports and are advertised independently for each port using this Sub-TLV. This per-port advertisement allows remote peers to construct appropriate NAT traversal parameters for each potential tunnel endpoint.

Unlike pairwise NAT traversal mechanisms such as IKEv2 [RFC7296], which require peers to dynamically discover NAT properties during tunnel setup, the BGP-controlled SD-WAN architecture enables each SD-WAN edge node to proactively advertise its NAT properties to all peers through BGP signaling. This approach simplifies NAT traversal in large-scale SD-WAN deployments where each edge node may need to establish tunnels with many peers.

Sub-TLV Name: Extended Port Attribute

Sub-TLV Code: 65 (IANA assigned)

Sub-TLV Encoding: The encoding is shown in the figure below:

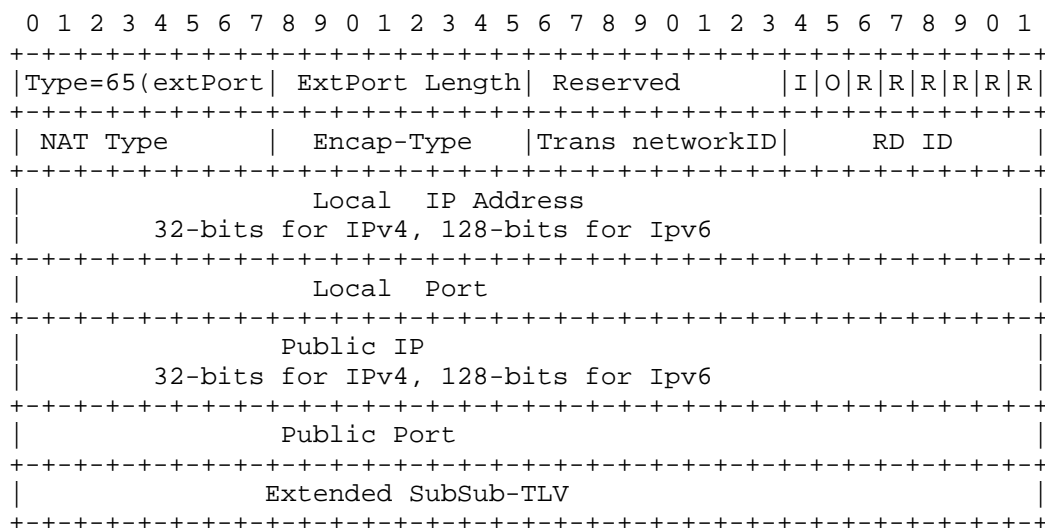


Figure 15: Extended Port Attribute Sub-TLV

where:

- * ExtPort Length: the length of the value field in octets excluding the Type and the Length fields. If IPv4, the length is 32 (8 header, 32 address, 8 for 1 SubSub-TLV). If IPv6, the length is 64 (8 header, 48 addresses, 8 for 1 subSubTLV).
- * Flags (16 bits):
 - I bit (C-PE port address or Inner address scheme):
 - o If set to 0, indicate the inner (private) address is IPv4.
 - o If set to 1, indicates the inner address is IPv6.
 - O bit (Outer address scheme):
 - o If set to 0, indicate the inner (private) address is IPv4.
 - o If set to 1, indicates the inner address is IPv6.
 - R bits: reserved for future use. MUST be set to 0, and ignored upon reception.

- * NAT Type (8 bites): an unsigned integer indicating the NAT behavior observed for this WAN port. The values are derived from the legacy NAT classification model described in RFC 3489 Section 5. The assigned values are:

- 1: without NAT ;
- 2: 1-to-1 static NAT;
- 3: Full Cone;
- 4: Restricted Cone;
- 5: Port Restricted Cone;
- 6: Symmetric; or
- 7: Unknown (i.e. no response from the STUN server).

The NAT Type value is determined by the sender using NAT discovery procedures (e.g., STUN [RFC8489] with legacy tests [RFC3489]) or local administrative configuration. The receiver is not required to verify NAT behavior but MUST validate that the received NAT Type field is within the range 1-7. Values outside this range are considered invalid and result in the Sub-TLV being treated as malformed.

- * Encap-Type(8 bits): An unsigned integer indicating the encapsulation type supported for this WAN port. This field is distinct from the Tunnel Type field in the BGP Tunnel Encapsulation Attribute [RFC9012]. The encapsulation types defined by this document are:

- Encap-Type=1: GRE;
- Encap-Type=2: VxLAN;

Notes:

- Other values are reserved for future specifications. The Encap-Type identifies the encapsulation protocol used within the IPsec payload when IPsec-SA Sub-TLVs (IPsec-SA-ID, IPsec-SA Nonce, IPsec Public Key, IPsec-SA Proposal, or Simplified IPsec-SA) are present in the SD-WAN Hybrid Tunnel.
- The Extended Port Attribute Sub-TLV does not support NAT traversal scenarios involving IPv4/IPv6 translation (e.g., NAT64 or 6to4).

- * Trans NetworkID (Transport Network ID) (8 bits): An identifier assigned by the SD-WAN Controller to indicate the transport network that this WAN port belongs to. All values from 0 to 255 are valid.
- * RD ID: The Routing Domain ID is a globally unique identifier assigned to the routing domain associated with this WAN port. All values from 0 to 255 are valid.
 - Some SD-WAN deployments may define multiple levels, zones, or regions that are represented as logical domains. Operational policies may govern whether tunnels are allowed between nodes in different logical domains. For example, a hub node may be permitted to establish tunnels across domains, while spoke nodes may be restricted to communicating only within their own domain. The definition, distribution, and enforcement of such policies are outside the scope of this document.
- * Local IP: The local (or private) IP address of the WAN port.
- * Local Port: used by Remote SD-WAN edge node for establishing IPsec to this specific port.
- * Public IP: The IP address after the NAT. If NAT is not used, this field is set to all-zeros
- * Public Port: The Port after the NAT. If NAT is not used, this field is set to all-zeros.
- * If NAT is not used for the WAN port, both the Public IP and Public Port fields MUST be set to zero. If one field is set to zero and the other is non-zero, the Sub-TLV is considered malformed.
- * Extended SubSub-TLV: for carrying additional information about the underlay networks.

If the Extended Port Attribute Sub-TLV is malformed (e.g., incorrect length, invalid address format, or unrecognized NAT type), it MUST be ignored per the procedures described in [RFC9012]. Other Sub-TLVs in the same Tunnel Encapsulation Attribute, if valid, MUST still be processed.

3.3.6.1. Extended Port SubSub-TLV

One Extended SubSub-TLVs is specified in this document: Underlay Network Type SubSub-TLV.

The Underlay Network Type SubSub-TLV is an optional SubSub-TLV used to advertise additional transport characteristics for the WAN port, including connection type, physical port type, and port bandwidth (e.g., LTE, DSL, Ethernet, and others). This information assists remote peers or controllers in selecting optimal underlay paths when multiple WAN ports are available. The Underlay Network Type SubSub-TLV is only valid for the Tunnel Type Hybrid SD-WAN within the Extended Port Attribute Sub-TLV.

Underlay Network Type.

66 (IANA Assigned).

The encoding is shown in the figure below:

```

  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| UnderlayType |           Length           |           Reserved           |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Connection Type | Port Type           |           Port Speed           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 16: Underlay Network Type SubSub-TLV

Where:

UnderlayType: Underlay Network Type (66 assigned by IANA)

Length: always 6 bytes

Reserved: 2-octet of reserved bits. It SHOULD be set to zero on transmission and MUST be ignored on receipt.

Connection Type: An unsigned integer indicating the connection type for this WAN port. Only a single value is carried per instance. The following values are defined:

- * 1 = Wired
- * 2 = WIFI
- * 3 = LTE
- * 4 = 5G
- * Values outside the range 1-4 are invalid and render the Sub-TLV malformed.

Port Type: An unsigned integer indicating the physical port type of the WAN interface. Only a single value is carried per instance. The following values are defined:

- * 1 = Ethernet
- * 2 = Fiber Cable
- * 3 = Coax Cable
- * 4 = Cellular
- * Values outside the range 1-4 are invalid and render the Sub-TLV malformed.

Port Speed: An unsigned 16-bit integer representing the port speed in megabits per second (Mbps). For example, a value of 1000 represents a port speed of 1000 Mbps (1 Gbps). The valid range is 1-65535. A Port Speed value of 0 is invalid and renders the Sub-TLV malformed.

Underlay Network Type SubSub-TLV is a MALFORMED SubSub-TLV if the fields do not fit the limits specified above. If a MALFORMED SubSub-TLV is contained in the Extended Port Attribute Sub-TLV, then the Extended Port Attribute Sub-TLV is MALFORMED. Per [RFC9012] a MALFORMED Sub-TLV is ignored.

3.4. Procedure for Client Routes with Hybrid SD-WAN Tunnel

Client routes with NLRI of AFI/SAFI IPv4 Unicast (1/1), IPv6 (2/1), L3VPN v4 Unicast (1/128), and IPv6 L3VPN (2/128) that use the Hybrid SD-WAN Tunnel Type can be advertised using one of two mechanisms:

Encapsulation Extended Community (Encap-EC) with SD-WAN SAFI: In this approach, the client route is advertised using Encapsulation Extended Community, as defined as "Barebones" in [RFC9012], to indicate the SD-WAN hybrid tunnel type. The detailed tunnel properties, such as IPsec-SAs, WAN port attributes, NAT properties, and other parameters, are advertised separately via BGP UPDATES using the SD-WAN SAFI. The SD-WAN Node ID, carried as the NextHop in client route advertisements and as the SD-WAN Node ID in SD-WAN SAFI underlay route advertisements, enables receiving BGP nodes to associate client routes with the correct underlay tunnels.

Tunnel Encapsulation Attribute (TEA): Alternatively, client routes

UPDATE can include all tunnel-related information directly in the same BGP UPDATE using the Tunnel Encapsulation Attribute (TEA). This encompasses the Hybrid SD-WAN Tunnel TLV along with its associated sub-TLVs, for example, those specifying IPsec proposals, public keys, nonces, NAT properties, and WAN port attributes. The NextHop attribute identifies the originating SD-WAN node, while the Tunnel Egress Endpoint Sub-TLV specifies the exact WAN port that terminates the tunnel.

The TEA-based approach, which include all tunnel attributes within route advertisement, can simplify the processing at the receiving nodes. However, it may lead to significant BGP attribute overhead, particularly when multiple IPsec-SAs are eligible to carry the same client route. In contrast, the Encap-EC approach (the "barebones" method defined in [RFC9012]) combined with SD-WAN SAFI separates tunnel attributes from route updates, enhancing flexibility and allowing tunnel properties to be reused across multiple client routes.

The SD-WAN Secure Links topology is supported using unicast IPv4 and IPv6 routes. L3VPN topologies, on the other hand, support the formation of Secure SD-WAN L3VPNs as described in [SD-WAN-BGP-USAGE] and MEF specifications [MEF 70.1] and [MEF 70.2].

3.4.1. SD-WAN Hybrid Tunnel Type in Encapsulation Extended Community

When client routes are advertised using the Encapsulation Extended Community (Encap-EC) with the SD-WAN Hybrid Tunnel Type, as specified in [RFC9012], the Encap-EC identifies the tunnel type, and the NextHop field in the BGP UPDATE serves as the Tunnel Egress Endpoint. Validation of the Tunnel Egress Endpoint follows the procedures defined in Sections 6, 8, and 13 of [RFC9012], as applied to the NextHop.

The Color Extended Community (Color-EC) is used to associate a client route with its eligible underlay tunnels. The Color value in the client route identifies the set of underlay tunnels, previously advertised with the same Color via SD-WAN SAFI, that may be used to transport the traffic. This enables SD-WAN ingress nodes or controllers to apply path selection policies based on performance, cost, or service requirements.

In this approach, if a required underlay tunnel is unavailable, the associated route MUST NOT be installed in the forwarding table or used to forward traffic. The route MAY still exist in the BGP control plane but MUST be marked as unusable for forwarding until a valid secure tunnel is established.

3.4.2. SD-WAN Hybrid Type in Tunnel Attributes via TEA

When client routes are advertised using the Tunnel Encapsulation Attribute (TEA) with the SD-WAN Hybrid Tunnel Type, the following procedures apply for validating the BGP UPDATE message:

1. Check for Well-formed SD-WAN Hybrid Tunnel TLV: A well-formed SD-WAN Hybrid Tunnel TLV MUST include a Tunnel

Egress Endpoint Sub-TLV if the hybrid tunnel terminates at a specific WAN port. If the tunnel is intended to terminate at the SD-WAN node level, the Tunnel Egress Endpoint Sub-TLV MAY be omitted. The validation for the Tunnel Egress Endpoint uses the [RFC9012] validation in section 6, 8, and 13. An invalid Tunnel Egress Endpoint, cause the Hybrid SD-WAN Tunnel TLV to be invalid, and the TLV is ignored.

It MAY also have any of the following Sub-TLVs:

- * the Color Sub-TLV defined in [RFC9012],
- * IPsec-SA-ID,
- * IPsec-SA Rekey Counter,
- * IPsec Public Key,
- * IPSec SA Proposal, or
- * Simplified IPsec-SA

A MALFORMED Sub-TLV is ignored in the Tunnel TLV is ignored.

Sub-TLV with an unknown type is ignored.

2. Check for multiple instances of Sub-TLVs: As specified in [RFC9012], only the first instance of a Sub-TLV is processed; subsequent ones are ignored.

A SD-WAN Hybrid tunnel TLV MAY have multiple instances of the IPsec-SA-ID if the IPsec-SA Identifiers are unique. If all the IPsec-SA Identifies are not unique, the second Sub-TLV is ignored and not propagated.

3. Validate Tunnel Egress Endpoint: The Tunnel Egress Endpoint MUST link to the remote end of one of the underlay links to be used. This validation adheres to the [RFC9012] Tunnel Egress Endpoint validation. The tunnel link MAY be active or inactive.
4. Validate each NLRI: Local policy is run to validate routes.

5. Validate Next Hop: The next hop MUST be be reachable via the tunnel."

3.4.3. Client Routes Carried Over Multiple Hybrid SD-WAN Tunnels

When a client route is advertised with the Encapsulation Extended Community (Encap-EC) that identifies the Hybrid SD-WAN Tunnel Type, the route may also include a Color Extended Community (Color-EC). This combination allows the route to be carried over multiple underlay tunnels that were previously advertised, each with the same Color value.

The Color-EC serves as a correlation mechanism: all underlay tunnels that have been advertised (via SD-WAN SAFI) with the same Color value are considered eligible to carry the traffic for the client route. This approach supports flexible path selection and tunnel diversity while avoiding the need to enumerate each tunnel per route.

This model is especially useful when:

- * A site has multiple available IPsec tunnels or WAN links.
- * A centralized controller or ingress SD-WAN edge node must select the optimal tunnel for forwarding based on performance, policy, or service constraints.

The tunnel attributes, including IPsec parameters, NAT traversal info, and WAN port properties, are conveyed separately via SD-WAN SAFI updates. This keeps client route updates minimal, allowing multiple routes to reference the same tunnel attributes by using the Color-EC.

3.4.4. SD-WAN VPN ID in Control Plane

In a BGP-controlled SD-WAN network, the VPN ID distinguishes client VPNs and ensures route separation. It is conveyed in client route UPDATES as follows:

- * For IPv4/IPv6 Unicast (AFI/SAFI = 1/1 or 2/1), the Route Target Extended Community SHOULD be included. The Route Target value is interpreted as the VPN ID. The Route Target is especially necessary when the SD-WAN edge node serves multiple VPNs on its client-facing interfaces. If all client routes belong to a single VPN and the association is unambiguous, the Route Target MAY be omitted.
- * For VPN-IPv4/VPN-IPv6 (AFI/SAFI = 1/128 or 2/128), the RD in the NLRI serves as the VPN ID.

3.4.5. SD-WAN VPN ID in Data Plane

In the data plane, SD-WAN traffic can traverse either an MPLS or IPsec segment within a Hybrid SD-WAN tunnel. The method for conveying the VPN ID depends on the encapsulation:

- * MPLS Segments: When the Hybrid tunnel uses MPLS transport, the MPLS label stack is used to identify the VPN per [RFC8277]. Security is assumed to be provided by the MPLS transport.
- * IPsec Segments: When traversing a public network with IPsec encryption: For GRE encapsulation within IPsec, the GRE Key field can carry the SD-WAN VPN ID; For VXLAN network virtualization overlays within IPsec, the VNI (Virtual Network Identifier) field is used to carry the VPN ID.

3.5. Procedure for Underlay Routes with Hybrid SD-WAN Tunnel TLV

Underlay routes in a BGP-controlled SD-WAN network are advertised using the SD-WAN SAFI, with the Tunnel Encapsulation Attribute (TEA) carrying a Hybrid SD-WAN Tunnel TLV. This TLV includes one or more Sub-TLVs that describe detailed tunnel attributes of the SD-WAN edge node's WAN ports, such as encapsulation types, NAT behavior, bandwidth, and IPsec parameters.

These underlay route advertisements carry the tunnel attributes needed for establishing SD-WAN Hybrid tunnels. Remote nodes use the SD-WAN Node ID carried in the SD-WAN SAFI to correlate client routes whose NextHop address matches the Node ID. This allows the receiving node to associate each client route with the appropriate set of tunnel attributes advertised by the corresponding SD-WAN edge node.

3.5.1. Hybrid SD-WAN NLRI without Encapsulation Extended Community

The Hybrid SD-WAN NLRI MUST be accompanied with the TEA, and MUST NOT be accompanied by an Encapsulation Extended Community.

3.5.2. Underlay Route with a TEA

The procedures for processing underlay routes follows the following steps:

1. Check for Well-Formed SD-WAN Hybrid Tunnel TLV: An Hybrid SD-WAN Tunnel TLV is well-formed using only Sub-TLVs valid for association with the underlay Route.

A well-formed SD-WAN Hybrid Tunnel TLV includes tunnel

attributes associated with a specific SD-WAN edge node, identified by the SD-WAN Node ID. The presence of the Tunnel Egress Endpoint sub-TLV indicates that the tunnel terminates at a specific WAN port on the SD-WAN node. If this sub-TLV is absent, the tunnel is considered to terminate at the node level, allowing any of the node's WAN ports to be used.

The Hybrid SD-WAN NLRI MUST NOT be accompanied by an Encapsulation Extended Community.

The SD-WAN Hybrid Tunnel TLV may contain the following sub-TLVs: Tunnel Egress Endpoint, IPsec-SA-ID, IPsec-SA Rekey Counter, IPsec Public Key, IPsec-SA Proposal, Simplified IPsec-SA, and Extended Port Attribute.

Following [RFC9012] intent, a MALFORMED Sub-TLV is ignored, and a sub-TLV with an unknown type is ignored.

2. Multiple instances of Sub-TLVs within a SD-WAN Tunnel TLV: As specified in [RFC9012], only the first instance of a Sub-TLV is processed; subsequent ones are ignored. The IPsec-SA-ID sub-TLVs MAY have multiple instances of the sub-TLV if the IPsec-SA Identifiers are unique, but if the IPsec-SA Identifiers are not unique the second sub-TLV is ignored and not propagated. If multiple Extended Port Sub-TLVs exist, the TLVs must be validated in step 4.
3. Validate Tunnel Egress Endpoint: The Tunnel Egress Endpoint MUST link to remote end point of one of the underlay links from the router receiving the link.
4. Validate Extended Port Attribute Sub-TLV(s): As described in Section 4.3.6, each Extended Port Attribute sub-TLV describes the properties of a single WAN port. Therefore, multiple Extended Port sub-TLVs may be present when the SD-WAN edge node has multiple WAN ports. Each sub-TLV MUST be validated to ensure that the port information it contains is sufficient to support the establishment of a tunnel to the remote peer. If any Extended Port Attribute Sub-TLV is determined to be invalid, the entire SD-WAN Hybrid Tunnel TLV MUST be considered invalid.
5. Validate each NLRI: Each typed NLRI in the SD-WAN Underlay MUST be well-formed, meaning it conforms to the structure defined in Section 4.2.1, including correct field lengths and ordering. A MALFORMED NLRI MUST be discarded; implementations MAY log an error. For well-formed NLRIs, the route's acceptance MUST be determined by local policy, based on the contents (e.g., Node ID, Color).

6. Validate Next Hop: The IP address specified in the Next Hop field MUST be reachable by the Tunnels.

3.5.3. Underlay Routes with Port-Local-ID of Zero

As specified in Section 4.2.1, a Route Type 1 NLRI includes the tuple (Port-Local-ID, SD-WAN-Color, SD-WAN-Node-ID). The Port-Local-ID field MAY be set to zero to indicate that the NLRI applies to all WAN ports on the identified SD-WAN node, effectively representing tunnel attributes at the node level rather than a specific port.

When Port-Local-ID = 0, the receiving BGP speaker SHOULD apply local policy to determine how to associate client routes with underlay tunnels. This local policy may prefer tunnels from specific SD-WAN nodes, or choose among SD-WAN Colors based on administrative preference, link type, path performance, or service-level objectives. The exact selection logic is implementation-specific.

It is valid for multiple such node-level NLRIs to be received, each advertising different SD-WAN Colors for the same node. For example, the following three NLRIs may be received (within one or more UPDATE messages):

Port-Local-ID (0), SD-WAN-Color (10), SD-WAN-Node-ID (2.2.2.2),

Port-Local-ID (0), SD-WAN-Color (20), SD-WAN-Node-ID (2.2.2.2),
and

Port-Local-ID (0), SD-WAN-Color (30), SD-WAN-Node-ID (2.2.2.2).

These indicate that node 2.2.2.2 supports multiple tunnel groups, each classified by a different SD-WAN Color. For example, these Colors may correspond to service tiers such as gold, silver, and bronze. The SD-WAN-Color field is used to correlate underlay tunnels with client routes that carry a matching Color Extended Community. If no match is found, the client route may not be forwarded over any SD-WAN tunnel.

3.5.4. Multiple Tunnels attached to One Underlay Route

An underlay route (SD-WAN NLRI) MAY only attach to one Hybrid SD-WAN Tunnel. If there are more than one Hybrid SD-WAN Tunnel TLV within a single TEA, the first is processed and the subsequent Hybrid SD-WAN Tunnel TLVs are ignored.

3.6. Error handling

The Error handling for SD-WAN VPN support has two components: error handling for Tunnel Encapsulation signaling (Encapsulation Extended Community and TEA) and the SD-WAN NLRI. An SD-WAN NLRI, a Tunnel Encapsulation attribute MUST always accompany the SD-WAN NLRI.

The previous sections (3.4 and 3.5) provide the procedures for handling client routes and undelay routes.

3.6.1. Error handling for the Tunnel Encapsulation Signaling

The error handling for the tunnel encapsulation signaling (Encapsulation Extended Community and TEA) in this document follows the procedures specified in [RFC9012], Sections 6, 8, and 13. Unless otherwise stated, malformed or unrecognized Sub-TLVs MUST be handled as specified in [RFC9012]. This document defines new Sub-TLVs for Tunnel Type 25 (SD-WAN-Hybrid), but does not alter the validation behavior established in RFC 9012.

The Tunnel encapsulation signaled with the client routes indicates the Egress endpoint via Next Hop in the Encapsulation Extended Community or the TEA Sub-TLV for Tunnel Egress Endpoints. As indicated in the procedure in sections 3.4.2 and 3.5.2, the SD-WAN Hybrid tunnel follows the validation section 6, 8, and 13 from [RFC9012].

The SD-WAN client routes associate some of the NLRIs that [RFC9012] associates with the Encapsulation Extended Community and the TEA using the validation specified in [RFC9012] in sections 6, 8, and 13. When the SD-WAN Hybrid Tunnel is associated with the SD-WAN NLRI, and all RFC9012 validation rules in section 6, 8, and 13 are extended to apply to the SD-WAN NLRI.

[RFC9012] contains the necessary detail to specify validation for the new Sub-TLVs present for the SD-WAN Tunnel type. However, to aid users of this document the following recap of validation of [RFC9012] is provided below. The validation from section 13 of [RFC9012] includes:

- * Invalid tunnel type MUST be treated if the TLV was not present.
- * A malformed sub-TLVs MUST be handled per [RFC9012] Section 13. If Tunnel Egress Endpoint is malformed, the entire TLV MUST be ignored. For security-sensitive attributes, such as those related to IPsec-SA setup, malformed or invalid values MUST be discarded and MUST NOT be used in security association processing. The BGP UPDATE containing such attributes SHOULD still be processed if

other attributes remain valid. Implementations SHOULD log the error for operational awareness and MAY trigger a session reset or rekeying if required by local policy. Unlike general BGP attributes, failure to process security-related information correctly could lead to misconfigurations or weakened security.

- * Multiple incidents of Tunnel Egress Endpoint Sub-TLV cause the first incident of these sub-TLVs to be utilized. Subsequent TLVs after the first one per type are ignored (per RFC9012), but propagated.
- * If a sub-TLV is meaningless for a tunnel type, the sub-TLV is ignored, but the sub-TLV is not considered malformed or removed from the Tunnel Attribute propagated with the NLRI.

For SD-WAN client routes with a TEA with a SD-WAN Hybrid Tunnel type TLV, the IPsec Sub-TLVs (IPsec-SA-ID, IPsec nonce, IPsec Public Key, IPsec Proposal, and Simplified IPsec-SA) are meaningful, but MAY be rarely sent. Incorrect fields within any of these 5 TLVs. Per [RFC9012], a malformed sub-TLV is treated as an unrecognized sub-TLV.

For SD-WAN NLRI underlay routes, the Extended Port sub-TLV and the IPsec sub-TLVs (IPsec-SA-ID, IPsec nonce, IPsec Public Key, IPsec Proposal, and Simplified IPsec-SA) are valid and meaningful. Incorrect fields within any of these 6 TLVs or SubSub-TLVs within the TLVs SHOULD cause the sub-TLV to be treated as malformed sub-TLV. Per [RFC9012], a malformed sub-TLV is treated as an unrecognized sub-TLV.

If multiple instances of the IPsec nonce, IPsec Public Key, IPsec Proposal, and Simplified IPsec are received within a SD-WAN Tunnel TLV, only the first is processed. The second instance is ignored and not propagated. The IPsec-SA-ID MAY have multiple copies, but the IPsec-SA Identifiers sent in the second sub-TLV MUST be different than any in the first IPsec-SA ID sub-TLV.

If multiple instances of the Extended Port sub-TLV are received, the local policy MUST determine which is to be used.

3.6.2. Error Handling for NLRI

The SD-WAN NLRI [AFI 1/SAFI = 74] utilizes a route type field to describe the format of the NLRI. This specification only allows an NLRI with a type value of 1. An NLRI with a type of field of another value is ignored and not processed. The implementation MAY log an error upon the reception of a type value outside of Route Type 1. Error handling for the SD-WAN NLRI also adheres to the BGP UPDATE error handling specified in [RFC7606].

The local policy configuration in the BGP peer receiving this NLRI MUST determine the validity of the route based on policy. Local configuration and policy MUST carefully constrain the SD-WAN-NLRI, tunnels, and IPsec security associations to create a "walled garden".

3.6.3. SDWAN NLRI and Tunnel Encapsulation Attribute

The SD-WAN NLRI (AFI=1/SAFI=74) MUST be paired with Tunnel Encapsulation attribute with a tunnel TLV for tunnel type SD-WAN-Hybrid. If the SD-WAN NLRI exist in an BGP UPDATE without a Tunnel Encapsulation Attribute (TEA) with a tunnel TLV for tunnel type SD-WAN-Hybrid, the NLRI is considered malformed and Treat-as-withdraw approach (per RFC7606).

The SD-WAN NLRI SHOULD not be paired with an Encapsulation Extended Community. If an SD-WAN NLRI is paried with an Encapsulation Extended Community rather than a Tunnel Encapsulation Attribute, the SD-WAN NLRI is considered malformed and the Treat-as-withdraw approach (per [RFC7606]) SHOULD be used.

4. Operational Consistency and Tunnel Validation

Unlike MPLS VPN whose PE nodes are all controlled by the network operators, SD-WAN edge nodes can be installed anywhere, in shopping malls, in 3rd party Cloud DCs, etc.

It is essential to ensure that advertisements from an SD-WAN edge node are legitimate. The RR, which maintains policy information about which SD-WAN nodes are authorized to communicate, MUST verify that the advertising BGP speaker is permitted to originate SD-WAN Hybrid Tunnel information before reflecting such routes to other peers.

4.1. Detecting Misaligned Tunnels

It is critical that a Hybrid SD-WAN Tunnel forwards traffic in accordance with local policy, taking into account the client route attributes, tunnel ingress and egress endpoints, and the associated security parameters.

To maintain correctness and security, both the RR and BGP speakers SHOULD validate that the client routes and associated tunnel information are consistent with expected configurations. This includes verifying that:

- * The NextHop in the client route update matches a known SD-WAN Node ID.

- * The tunnel's egress endpoints are reachable and authorized.
- * The advertised SD-WAN Color in the underlay NLRI matches the Color Extended Community attached to the client route.

4.2. IPsec Attributes Mismatch

Each SD-WAN node (e.g., a C-PE) can advertise its IPsec-related attributes to remote peers using Sub-TLVs within the Tunnel Encapsulation Attribute, in one of the following three forms, to support the establishment of IPsec-SAs:

- * Identifiers of a pre-established IPsec-SA, carried in IPsec-SA-ID Sub-TLV.
- * a simplified set of security parameters for setting up a IPsec-SA, such as Transform type, IPsec Mode, AH/ESP Algorithms, rekey counter, 2 public keys, nonce, and duration, carried in the Simplified IPsec-SA Sub-TLV.
- * A flexible representation of IPsec parameters, where the Nonce, Public Key, and SA Proposal are individually specified and carried in the IPsec-SA Rekey Counter Sub-TLV, IPsec Public Key Sub-TLV, and IPsec-SA Proposal Sub-TLV, respectively.

For existing IPsec-SAs, an SD-WAN node that receives the advertisement can simply use one of the existing SAs to forward traffic for the associated client routes. If multiple SAs are available for a given client route, local policy on the receiving SD-WAN node MAY determine which SA is selected.

When a new IPsec-SA is to be established using parameters carried in Sub-TLVs, such as the IPsec-SA Rekey Counter Sub-TLV, IPsec Public Key Sub-TLV, and IPsec-SA Proposal Sub-TLV, the receiving SD-WAN node MUST validate that the proposed IPsec transforms and algorithms are compatible with its local configuration. These attributes, received via the Tunnel Encapsulation Attribute (TEA), define the parameters for establishing the IPsec tunnel between local and remote WAN ports. This compatibility check is performed at the IPsec layer, not by BGP.

The C-PE devices do not attempt to negotiate IPsec-SA parameters or transform sets with remote peers. Instead, the configurations must match as advertised. If there is a mismatch, either in the simple IPsec-SA identifiers or in the detailed transform parameters, no tunnel is established. Implementations MAY discard incompatible proposals or log them for operational visibility.

4.2.1. Example creation of IPsec-SA over SD-WAN Hybrid tunnel

This section provides an example of how IPsec-SA is established over an SD-WAN Hybrid tunnel. Figure 1 in Section 4 illustrates an IPsec tunnel being created between port P2 (192.168.1.10) on C-PE1 and port P2 (192.168.2.1) on C-PE2.

To establish this tunnel, C-PE1 must advertise the following attributes required for setting up the IPsec-SA:

- * NextHop: 192.168.1.10
- * SD-WAN Node ID: 1.1.1.1
- * SD-WAN-Site-ID: 15.0.0.2
- * Tunnel Encap Attr (Type=SD-WAN) -
 - Extended Port Attribute Sub-TLV containing
 - o Transport SubSub-TLV - with information on ISP.
 - IPsec information for detailed information about the ISP
 - IPsec-SA Rekey Counter Sub-TLV,
 - IPsec-SA Public Key Sub-TLV,
 - Proposal Sub-TLV (type = ENCR, transform ID = 1)
 - o type: ENCR
 - o Transform ID: 1
 - o Transform attributes = trans 1 [from RFC7296]

C-PE2 needs to advertise the following attributes for establishing IPsec-SA:

Next Hop: 192.168.2.1

SD-WAN Node ID: 2.2.2.2

SD-WAN-Site-ID: 1500

Tunnel Encap Attr (Type=SD-WAN)

- * Extended Port Attribute Sub-TLV

- Transport SubSub-TLV - with information on ISP.
- * IPsec-SA Rekey Counter Sub-TLV,
- * IPsec-SA Public Key Sub-TLV,
- * IPSec Proposal Sub-TLV with
 - transform type: ENCR
 - Transform ID = 1
 - Transform attributes = trans 2

As there is no matching transform between the WAN ports P2 and P2 in C-PE1 and C-PE2, respectively, no IPsec Tunnel will be established.

5. Manageability Considerations

The BGP-based signaling mechanisms described in this document are primarily intended to enable SD-WAN edge nodes to advertise underlay transport and tunnel parameters to their RR. These parameters, once received, can be monitored and validated using existing BGP monitoring tools such as BMP or route policy inspection frameworks. Operators SHOULD implement logging and alerting mechanisms for cases where inconsistent or malformed Sub-TLVs are received, as specified in Section 4.6. Misaligned parameters, such as mismatched IPsec-SA-IDs or invalid NAT indicators, should trigger operational alerts to aid troubleshooting. No new MIB modules or YANG models are introduced in this document, but implementations are expected to expose relevant state (e.g., tunnel type, advertised properties) via standard operational interfaces. The use of secure transport connections (e.g., BGP over IPsec/TLS) is RECOMMENDED to ensure manageability in untrusted environments.

6. Security Considerations

This document defines BGP extensions for SD-WAN edge nodes to advertise their attributes for establishing IPsec-SAs and underlay tunnel attributes, typically via a RR, which then propagates them to authorized SD-WAN peers. These BGP messages may contain sensitive information such as public keys, IPsec proposals, and nonces. In deployments where SD-WAN edge nodes communicate with the RR over public or untrusted networks, BGP MUST be run over a secure transport, such as TCP protected by IPsec or TLS. These secure channels protect all fields, including cryptographic attributes, from tampering or interception. Without such protection, the system may be vulnerable to spoofed tunnel attributes, unauthorized route

injections, or replayed IPsec setup information.

However, in closed or "walled garden" deployments, where SD-WAN edge nodes and the RR (SD-WAN controller) are within a trusted, secured environment (e.g., a private MPLS backbone or physically secured enterprise network), the risk of interception or tampering is significantly reduced. In such cases, the use of secure transport is optional, and operators may choose to run BGP over standard TCP, based on their internal risk assessment.

Regardless of the transport used, BGP policy enforcement remains critical. The RR SHOULD apply strict filtering and policy controls to validate that only authorized SD-WAN edge nodes advertise specific Node IDs, Route Targets, or VPN identifiers. While route origin validation via RPKI helps, it does not cover SD-WAN-specific fields like Tunnel attributes or SA proposals. Local policies, when misconfigured, may introduce vulnerabilities; therefore, policy application points SHOULD be carefully audited.

Many of the general BGP security risks discussed here are also covered in [RFC4271], [RFC4272], and [RFC9012]. This document inherits those considerations and introduces no new cryptographic requirements beyond what is described for securing BGP transport and validating the correctness of SD-WAN tunnel attribute exchanges.

This specification does not define deployments across fully untrusted networks, but if such environments are used, strong transport security becomes a MUST, and additional validation mechanisms may be required to maintain SD-WAN tunnel and routing integrity.

7. IANA Considerations

7.1. SD-WAN Overlay SAFI

IANA has assigned SAFI = 74 as the SD-WAN SAFI.

7.2. Tunnel Encapsulation Attribute Tunnel Type

IANA is requested to assign a type from the BGP Tunnel Encapsulation Attribute Tunnel Types as follows [RFC8126]:

Value	Description	Reference
-----	-----	-----
25	SD-WAN-Hybrid	(this document)

7.3. Tunnel Encapsulation Attribute Sub-TLV Types

IANA is requested to assign the following sub-Types in the BGP Tunnel Encapsulation Attribute Sub-TLVs registry:

Value	Type Description	Reference	Section
64	IPSEC-SA-ID Sub-TLV	This document	4.3.1
65	Extended Port Attribute Sub-TLV	This document	4.3.6
66	Underlay Type SubSub-TLV	This document	4.3.6.1
67	IPsec-SA Rekey Counter Sub-TLV	This document	4.3.2
68	IPsec Public Key Sub-TLV	This document	4.3.3
69	IPsec-SA Proposal Sub-TLV	This document	4.3.4
70	Simplified IPsec-SA sub-TLV	This document	4.3.5

7.4. SD-WAN Edge Discovery NLRI Route Types

IANA is requested to create a new registry titled "SD-WAN Edge Discovery NLRI Route Types" under the "Border Gateway Protocol (BGP) Parameters" registry. The allocation policy for this registry shall be IETF Review (as defined in RFC 8126):

Value	Description	Reference
1	SD-WAN Tunnel Endpoint NLRI Route Type	(this document)

Values 2-255 are reserved for future assignments.

7.5. SD-WAN Extended Port Encapsulation Types

IANA is requested to create a new registry titled "SD-WAN Extended Port Encapsulation Types" under the BGP Tunnel Encapsulation Sub-TLV registries.

Value	Type Description	Reference
0	Reserved	This document
1	GRE	This document
2	VXLAN	This document
3~255	Reserved for future	

7.6. SD-WAN Extended Port Connection Types

IANA is requested to create a new registry titled "SD-WAN Extended Port Connection Types" under the BGP Tunnel Encapsulation Sub-TLV registries.

Value	Type Description	Reference
-----	-----	-----
0	Reserved	This document
1	Wired	This document
2	WIFI	This document
3	LTE	This document
4	5G	This document
5~255	Unassigned	
255	Reserved for Experimental Use	

7.7. SD-WAN Extended Port Physical Port Types

IANA is requested to create a new registry titled "SD-WAN Extended Port Physical Port Types" under the BGP Tunnel Encapsulation Sub-TLV registries.

Value	Type Description	Reference
-----	-----	-----
0	Reserved	This document
1	Ethernet	This document
2	Fiber Cable	This document
3	Coax Cable	This document
4	Cellular	This document
5~255	Unassigned	
255	Reserved for Experimental Use	

8. References

8.1. Normative References

- [MEF70.1] MEF, "SD-WAN Service Attributes and Service Framework", November 2021, <<https://www.mef.net/resources/mef-70-1-sd-wan-service-attributes-and-service-framework/>>.
- [MEF70.2] MEF, "SD-WAN Service Attributes and Service Framework", October 2023, <<https://www.mef.net/resources/mef-70-2-sd-wan-service-attributes-and-service-framework/>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<https://www.rfc-editor.org/info/rfc4456>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC5114] Lepinski, M. and S. Kent, "Additional Diffie-Hellman Groups for Use with IETF Standards", RFC 5114, DOI 10.17487/RFC5114, January 2008, <<https://www.rfc-editor.org/info/rfc5114>>.
- [RFC5903] Fu, D. and J. Solinas, "Elliptic Curve Groups modulo a Prime (ECP Groups) for IKE and IKEv2", RFC 5903, DOI 10.17487/RFC5903, June 2010, <<https://www.rfc-editor.org/info/rfc5903>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.
- [RFC8489] Petit-Huguenin, M., Salgueiro, G., Rosenberg, J., Wing, D., Mahy, R., and P. Matthews, "Session Traversal Utilities for NAT (STUN)", RFC 8489, DOI 10.17487/RFC8489, February 2020, <<https://www.rfc-editor.org/info/rfc8489>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.

8.2. Informative References

- [IANA-AH] IANA, "IANA-AH", <<https://www.iana.org/assignments/isakmp-registry/isakmp-registry.xhtml#isakmp-registry-9>>.
- [IANA-ESP] IANA, "IANA-ESP", <<https://www.iana.org/assignments/isakmp-registry/isakmp-registry.xhtml#isakmp-registry-9>>.
- [Net2Cloud]
L. Dunbar, A Malis, C. Jacquenet, M. Toy and K. Majumdar, "Dynamic Networks to Hybrid Cloud DCs: Problem Statement and Mitigation Practice", September 2023, <<https://datatracker.ietf.org/doc/draft-ietf-rtgwg-net2cloud-problem-statement/>>.
- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<https://www.rfc-editor.org/info/rfc1997>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, DOI 10.17487/RFC4659, September 2006, <<https://www.rfc-editor.org/info/rfc4659>>.

- [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, DOI 10.17487/RFC4761, January 2007, <<https://www.rfc-editor.org/info/rfc4761>>.
- [RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, DOI 10.17487/RFC4762, January 2007, <<https://www.rfc-editor.org/info/rfc4762>>.
- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", RFC 5701, DOI 10.17487/RFC5701, November 2009, <<https://www.rfc-editor.org/info/rfc5701>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <<https://www.rfc-editor.org/info/rfc6793>>.
- [RFC7018] Manral, V. and S. Hanna, "Auto-Discovery VPN Problem Statement and Requirements", RFC 7018, DOI 10.17487/RFC7018, September 2013, <<https://www.rfc-editor.org/info/rfc7018>>.
- [RFC8092] Heitz, J., Ed., Snijders, J., Ed., Patel, K., Bagdonas, I., and N. Hilliard, "BGP Large Communities Attribute", RFC 8092, DOI 10.17487/RFC8092, February 2017, <<https://www.rfc-editor.org/info/rfc8092>>.
- [SD-WAN-BGP-USAGE]
L. Dunbar, A Sajassi, J Drake, and B. Najem, "BGP Usage for SD-WAN Overlay Networks", September 2023, <<https://datatracker.ietf.org/doc/draft-ietf-bess-bgp-sdwan-usage/>>.
- [Secure-EVPN]
A Sajassi, A. Banerjee, S. Thoria, D. Carrell, B. Weis, J. Drake, "Secure EVPN", November 2024, <<https://datatracker.ietf.org/doc/draft-ietf-bess-secure-evpn/>>.

Appendix A. Acknowledgments

Acknowledgements to Wang Haibo, Shunwan Zhuang, Hao Weiguo, and ShengCheng for implementation contribution. Many thanks to Yoav Nir, Graham Bartlett, Jim Guichard, John Scudder, and Donald Eastlake for their review and suggestions.

Contributors

Below is a list of other contributing authors:

- * Gyan Mishra,
- * Shunwan Zhuang,
- * Sheng Cheng, and
- * Donald Eastlake.

Authors' Addresses

Linda Dunbar
Futurewei
Dallas, TX,
United States of America
Email: ldunbar@futurewei.com

Susan Hares
Huawei
United States of America
Email: shares@ndzh.com

Kausik Majumdar
Oracle
California,
United States of America
Email: kausik.majumdar@oracle.com

Robert Raszuk
Arrcus
United States of America
Email: robert@raszuk.net

Venkit Kasiviswanathan
Arista
United States of America
Email: venkit@arista.com