

Network Working Group
Internet Draft
Intended status: Informational
Expires: June 27, 2025

L. Dunbar
Futurewei
A. Sajassi
Cisco
J. Drake
Independent
B. Najem
Bell Canada
S. Hares
May 27, 2026

BGP Usage for SD-WAN Overlay Networks
draft-ietf-bess-bgp-sdwan-usage-36

Abstract

This document illustrates how a BGP-based control plane can be used to manage large scale Software Defined WAN (SD-WAN) overlay networks by distributing edge service reachability information, WAN port attributes, and underlay path details, thereby minimizing manual provisioning. In such deployments, BGP can provide a standards-based mechanism for distributing information that may otherwise be exchanged using proprietary SD-WAN control-plane mechanisms.

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."

Copyright Notice

Copyright (c) 2026 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction.....	3
2. Conventions used in this document.....	5
3. SD-WAN Scenarios and Their Requirements.....	6
3.1. SD-WAN Functional Overview.....	7
3.1.1. Supporting SD-WAN Segmentation.....	7
3.1.2. Client Service Requirement.....	7
3.1.3. SD-WAN Traffic Segmentation.....	8
3.1.4. Zero Touch Provisioning.....	9
3.1.5. Constrained Propagation of SD-WAN Edge Properties....	9
3.2. Scenario #1: Homogeneous Encrypted SD-WAN.....	10
3.3. Scenario #2: Differential Encrypted SD-WAN.....	11
3.4. Scenario #3: Private VPN PE based SD-WAN.....	13
4. Provisioning Model.....	14
4.1. Client Service Provisioning Model.....	14
4.2. Policy Configuration.....	14
4.3. IPsec Related Parameters Provisioning.....	15
5. BGP Controlled SD-WAN.....	15
5.1. Rationale for Using BGP as Control Plane for SD-WAN.....	15
5.2. BGP Scenario for Homogeneous Encrypted SD-WAN.....	17
5.3. BGP Scenario for Differential Encrypted SD-WAN.....	17
5.4. BGP Scenario for Flow-Based Segmentation.....	19
6. SD-WAN Forwarding Model.....	20
6.1. Forwarding Model for Homogeneous Encrypted SD-WAN.....	21
6.1.1. Network and Service Startup Procedures.....	21
6.1.2. Packet Walk-Through.....	21
6.2. Forwarding Model for Hybrid Underlay SD-WAN.....	22
6.2.1. Network and Service Startup Procedures.....	22
6.2.2. Packet Walk-Through.....	22

6.3. Forwarding Model for PE based SD-WAN.....	24
6.3.1. Network and Service Startup Procedures.....	24
6.3.2. Packet Walk-Through.....	24
7. Manageability Considerations.....	25
8. Security Considerations.....	26
9. IANA Considerations.....	28
10. References.....	28
10.1. Normative References.....	28
10.2. Informative References.....	29
11. Acknowledgments.....	31

1. Introduction

Software Defined Wide Area Network (SD-WAN), as described in [MEF70.2], provides overlay connectivity services that transport IP packets across one or more underlay networks by identifying traffic types and applying policies to determine forwarding behavior. Key characteristics of SD-WAN networks include:

- Transport Augmentation: an SD-WAN path can utilize different types of underlay networks, including private networks (with or without encryption) and public networks (commonly requiring encryption).
- Direct Internet Breakout: Traffic from remote branch offices can directly access the internet, avoiding backhauling to corporate headquarters for centralized policy control.
- Policy-Based Traffic Steering: Traffic can be directed over specific overlay paths based on predefined conditions, such as matching one or multiple fields in the IP header, rather than solely relying on destination IP addresses [RFC9522]. For IPv6 [RFC8200], attributes like the Flow Label, source address, specific extension header fields, or a combination of these can be used. Additional details are available in Tables 8 and 9 of [MEF70.2].
- Performance-Based Forwarding: Traffic can be steered based on performance metrics (e.g., packet delay, loss, jitter), selecting the underlay path that meets or exceeds policy requirements.

This document focuses on the use of BGP as the SD-WAN overlay control plane. While [Net2Cloud-Problem] identifies general issues encountered when interconnecting networks to cloud environments, this document discusses how BGP can be used to distribute

reachability, policy-related, and underlay path information in SD-WAN deployments. Additional operational drivers for standardized protocol behavior are summarized in Section 6 of [MPLIFY-119].

It's important to distinguish the BGP instance as the control plane for SD-WAN overlay from the BGP instances governing the underlay networks. The document requires that communication between the SD-WAN controller and SD-WAN edges for exchanging control-plane information be carried over a secure channel. In this document, references to "the Route Reflector (RR)" denote the RR instance(s) designated by the SD-WAN controller to distribute SD-WAN routing information. The detailed BGP extensions used for SD-WAN edge discovery and attribute distribution are specified in [SD-WAN-Discovery]. Deployments may integrate the controller and RR or keep them as separate components; the SD-WAN control plane is logically centralized but may be physically realized by a set of distributed controller/RR instances operating as a coordinated system for scalability, resiliency, and geographic redundancy, with the RR remaining the logical control point for SD-WAN route distribution and policy enforcement as described in this document; the mechanisms for coordination and state synchronization among multiple RR instances are out of the scope of this document.

The document assumes a single administrative domain for the SD-WAN overlay, consistent with the MEF SD-WAN service model. The underlay connectivity may be provided by one or more networks or NSPs, but the SD-WAN overlay control plane described in this document is under a single administrative control.

This document captures the SD-WAN scenarios described in Sections 3.2, 3.3, and 3.4, along with control-plane behaviors and forwarding considerations when BGP is used as the SD-WAN overlay control plane. Publishing this material as an RFC establishes a shared understanding of the SD-WAN problem space and deployment assumptions to assist with protocol development for BGP-based SD-WAN networks [SD-WAN-Discovery]. This document is informational and does not specify a complete SD-WAN solution. Its scope is limited to selected SD-WAN overlay control-plane functions, such as distributing edge reachability, WAN-port attributes, tunnel-related information, and policy-constrained routes.

2. Conventions used in this document

Cloud DC: Third party data centers that host applications and workloads owned by different organizations or tenants.

Controller: Used interchangeably with SD-WAN controller to denote the logical system that manages the SD-WAN overlay network. In a BGP-controlled SD-WAN, the controller may include a Route Reflector (RR) function for BGP route propagation and policy-based distribution of SD-WAN routing information. The RR function may be collocated with, embedded in, or otherwise integrated with the SD-WAN controller, but it represents only one component of the overall controller.

Client service: A service attached to the client-facing interface of an SD-WAN edge, identified by associated reachability or attachment information, such as an IP prefix or VLAN.

Client route: A BGP-advertised route originated by an SD-WAN edge that represents the reachability of a client-facing service (e.g., IP prefix or VLAN) and includes associated path attributes used by the SD-WAN Controller for policy enforcement and forwarding decisions.

C-PE: SD-WAN edge, which can be Customer Premises Equipment (CPE) for customer-managed SD-WAN, or Provider Edge (PE) for provider-managed SD-WAN services.

Homogeneous Encrypted SD-WAN: An SD-WAN network in which all traffic to/from the SD-WAN edges are carried by IPsec tunnels regardless of underlying networks. I.e., the client traffic is carried by IPsec tunnel even over MPLS private networks.

NSP: Network Service Provider.

PE: Provider Edge

- SD-WAN edge: A functional entity, either physical or virtual, that participates in the SD-WAN overlay network. These nodes advertise client routes to the SD-WAN Controller (e.g., BGP RR).
- SD-WAN: An overlay connectivity service that transport IP packets over one or more Underlay connectivity services by forwarding them based on policies. [MEF-70.2].
- SD-WAN IPsec SA: IPsec Security Association [RFC4301] between two WAN ports of the SD-WAN edges or between two SD-WAN edges.
- SD-WAN over Hybrid Underlay Networks: SD-WAN over Hybrid Underlay Networks typically have edge nodes utilizing bandwidth resources from different types of underlay networks, some being private networks and others being public Internet.
- WAN Port: A Port or Interface facing a Network Service Provider (NSP), with an address allocated by the NSP.
- Private VPN: refers to a VPN that is supported wholly by a single network service provider without using any elements of the public Internet and without any traffic passing out of the immediate control of that service provider.
- Zero Touch Provisioning (ZTP): a network automation approach that enables automatic provisioning and configuration of SD-WAN devices, such as routers and switches, at remote locations without manual intervention.

3. SD-WAN Scenarios and Their Requirements

This section outlines the core requirements for SD-WAN overlay networks and introduces various SD-WAN scenarios. Sections 3.2, 3.3, and 3.4 describe potential SD-WAN deployment scenarios, which are further explored in subsequent sections to illustrate how the BGP control plane can be used to distribute reachability and policy information within SD-WAN overlay networks.

3.1. SD-WAN Functional Overview

3.1.1. Supporting SD-WAN Segmentation

"SD-WAN Segmentation" refers to policy-driven network partitioning, a common approach in SD-WAN deployment. An SD-WAN segment is essentially a virtual private network (SD-WAN VPN) consisting of a set of edge nodes interconnected by tunnels, such as IPsec tunnels and/or MPLS VPN tunnels.

This document assumes that SD-WAN VPN configuration on PE devices will, as with MPLS VPN, make use of (Virtual Routing and Forwarding (VRF) instances [RFC4364] [RFC4659]. Notably, a single SD-WAN VPN can be mapped to one or multiple virtual topologies governed by the SD-WAN controller's policies.

When BGP is used for SD-WAN, the client route UPDATE is the same as MPLS VPN. The Route Target in the BGP Extended Community [RFC4360] can differentiate the routes belonging to different SD-WAN VPNs.

As SD-WAN is an overlay network arching over multiple types of networks, MPLS L2VPN [RFC4761] [RFC4762]/L3VPN [RFC4364] [RFC4659] or pure L2 underlay can continue using the VPN ID (Virtual Private Network Identifier), VN-ID (Virtual Network Identifier), or VLAN (Virtual LAN) in the data plane to differentiate packets belonging to different SD-WAN VPNs. For packets transported through an IPsec tunnel, additional encapsulation, such as GRE [RFC2784] or VxLAN [RFC7348], is needed to embed the SD-WAN VPN identifier inside the IPsec ESP header.

Section 3.1.3 further elaborates on traffic segmentation by providing operational examples and detailing how segmentation policies apply to individual flows; the two sections describe the same concept at different levels of granularity.

3.1.2. Client Service Requirement

The client service requirements describe the SD-WAN edge's ports, also known as SD-WAN client interfaces, which connect the client network to the SD-WAN service.

The SD-WAN client interface should support IPv4 and IPv6 addresses as well as Ethernet in accordance with the [IEEE802.3] standard.

In [MEF 70.2], the "SD-WAN client interface" is called SD-WAN UNI (User Network Interface). Section 4 of [MEF 70.2] defines a comprehensive set of attributes for the SD-WAN UNI, detailing the

expected behavior and requirements to enable seamless connectivity to the client network.

The client service at the SD-WAN edge must support the SD-WAN UNI service attributes outlined in Section 4 of [MEF 70.2].

3.1.3. SD-WAN Traffic Segmentation

SD-WAN Traffic Segmentation allows traffic to be separated based on business priorities, security requirements, and operational needs. This ensures that different user groups or services can operate within distinct topologies or follow tailored policies to meet specific business and security objectives.

For example, in a retail environment, traffic from point-of-sales (PoS) systems may require a different topology that is separate from other traffic. The PoS traffic is routed exclusively to the payment processing entity at a central hub site, while other types of traffic can be routed among all branches or remote sites.

In the figure below, traffic from the PoS system follows a tree topology (denoted as "----" in the figure below), whereas other traffic can follow a multipoint-to-multipoint topology (denoted as "===").

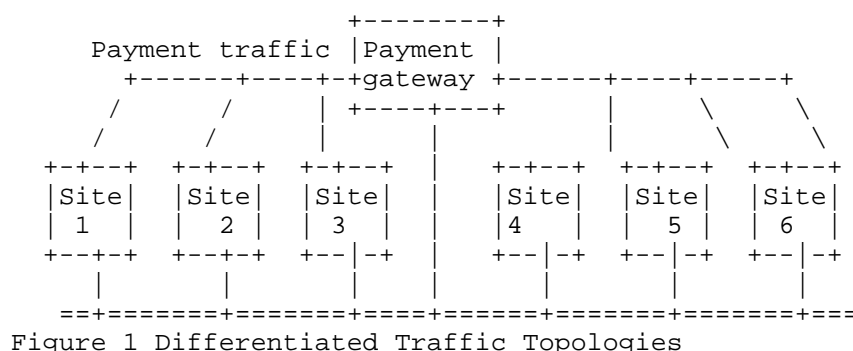


Figure 1 Differentiated Traffic Topologies

Another example is an enterprise that wants to isolate traffic by departments, ensuring each department having its own unique topology and policies. For instance, the HR department may need to access specific systems or resources that are not accessible by the engineering department. Similarly, contractors may have limited access to the enterprise resources.

3.1.4. Zero Touch Provisioning

SD-WAN Zero-Touch Provisioning (ZTP) allows an SD-WAN edge to obtain initial configuration without manual intervention. ZTP and onboarding procedures are not specific to BGP and are outside the scope of this document. In the context of this document, the only ZTP-related aspect is that the SD-WAN edge obtains the information needed to authenticate to the SD-WAN controller/RR and establish a secure channel for exchanging BGP UPDATE messages.

3.1.5. Constrained Propagation of SD-WAN Edge Properties

For an IPsec tunnel to be established between two edges for data exchange, both edges need to know each other's network properties, such as the IP addresses of the WAN ports, the edges' loopback addresses, the attached client routes, the supported encryption methods, etc., which are learned via BGP UPDATES exchanged through the RR.

In many cases, an SD-WAN edge is authorized to communicate with only a subset of other edge nodes. To maintain security and privacy, the property of an SD-WAN edge must not be propagated to unauthorized peers. However, when a remote SD-WAN edge powers up, it may lack the policies to determine which peers are authorized to communicate. Therefore, SD-WAN deployment needs to have a central point to distribute the properties of an SD-WAN edge to its authorized peers.

BGP is well suited for this purpose. A Route-Reflector (RR) [RFC4456], integrated into the SD-WAN controller, enforces policies governing the communication among SD-WAN edges. The RR ensures that BGP UPDATE messages from an SD-WAN edge are propagated only to other edges within the same SD-WAN VPN.

An SD-WAN edge must exchange BGP UPDATE messages with its designated RR over a protected transport. For example, IPsec [RFC4301] can be used to protect the path between the SD-WAN edge and RR, with the BGP TCP session established over that protected path. Other mechanisms for protecting the edge-to-RR BGP session are outside the scope of this document.




Figure 2: Authorized Peer Groups managed by RR

Tenant separation is achieved by the SD-WAN VPN identifiers represented in the control plane and data plane, respectively.

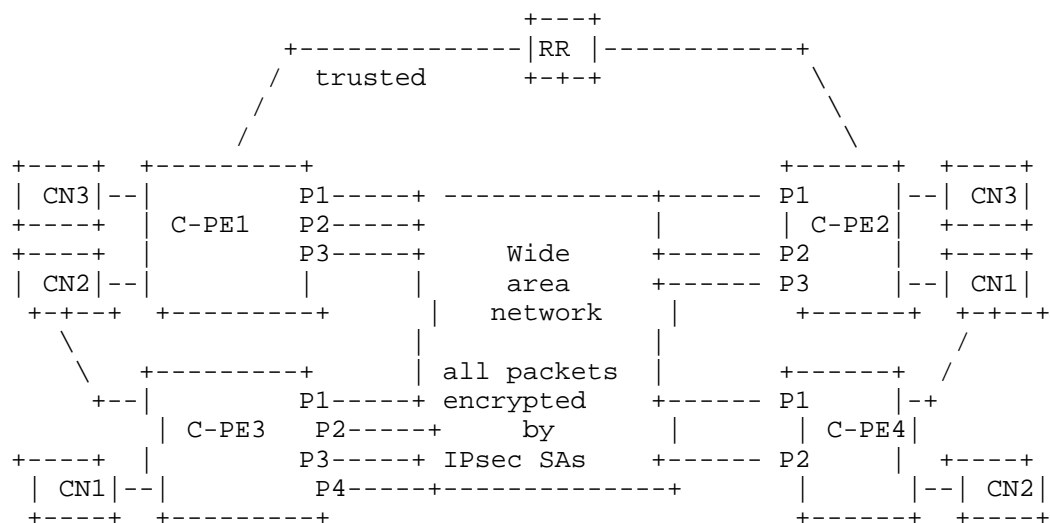
3.2. Scenario #1: Homogeneous Encrypted SD-WAN

Homogeneous Encrypted SD-WAN refers to an SD-WAN network where edge nodes encrypt all client traffic destined to other edge nodes, regardless of whether the underlay is private or public.

Typical use cases for Homogeneous Encryption:

- A small branch office connecting to its headquarters via the Internet. All traffic to and from this small branch office must be encrypted, usually achieved by IPsec Tunnels [RFC4301].
- A retail store in a shopping mall may need to securely connect to its services hosted in one or more Cloud DCs via the Internet. A common method involves establishing IPsec SAs with the Cloud DC gateway to securely transport sensitive data to/from the store.

The granularity of the IPsec SAs for Homogeneous Encryption can be per site, per subnet, per tenant, or per address. Once the IPsec SA is established for a specific subnet/tenant/site, all traffic to/from the subnet/tenant/site is encrypted.



CN: Client Networks, which is same as Tenant Networks used by NVo3
Figure 3: Homogeneous Encrypted SD-WAN

A Homogeneous Encrypted SD-WAN uses IPsec tunnels to protect traffic between SD-WAN edges. The SD-WAN overlay may include many edge nodes, with configuration and policy information managed by an SD-WAN controller.

Existing private VPNs (e.g., MPLS based) can use Homogeneous Encrypted SD-WAN to extend over the public network to remote sites to which the VPN operator does not own or lease infrastructural connectivity.

3.3. Scenario #2: Differential Encrypted SD-WAN

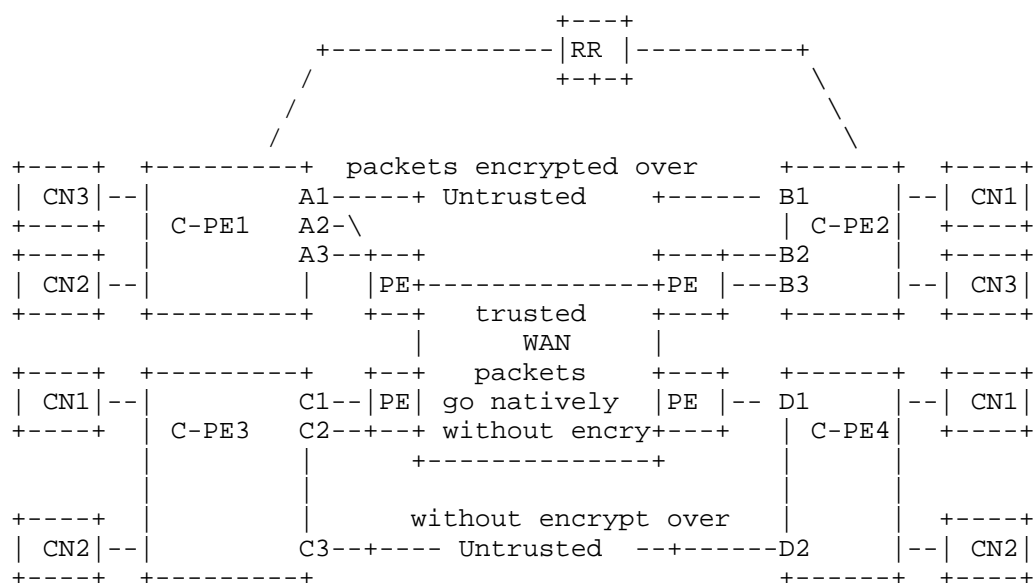
Differential Encrypted SD-WAN refers to an SD-WAN network that utilizes hybrid underlays, combining private VPNs and the public Internet. In this model, traffic traversing the private VPN is forwarded natively without encryption, while traffic over the public Internet is encrypted for security. This approach balances performance and security. Since IPsec encryption requires significant processing power and traffic over the public Internet typically lacks the premium SLA (Service Level Agreement) provided by private VPNs-especially over long distances-current practice is to forward traffic over private VPNs without encryption, leveraging the inherent reliability and security of the private network. Meanwhile, encryption is applied only to traffic routed over the public Internet to ensure data confidentiality.

One C-PE might have the Internet-facing WAN ports managed by different NSPs with the WAN ports' addresses assigned by the corresponding NSPs. Clients may define specific policies to govern how traffic flows across the network, such as:

- 1) Certain flows can only be forwarded over private VPNs.
- 2) Certain flows can be forwarded over either private VPNs or the public Internet. When forwarded over the public Internet, the packets are encrypted.
- 3) Some flows, especially Internet-bound browsing ones, can be handed off to the Internet without further encryption.

For example, consider a flow traversing multiple segments, A<->B, B<->C, C<->D, has Policy 2) above. This flow can cross different underlays in different segments, such as over Private underlay between A<->B without encryption or over the public Internet between B<->C protected by an IPsec SA.

In the figure below, each C-PE has both Internet-facing and private-VPN interfaces (e.g., A1, B2, C3, and D2 connect to the Internet, while others connect to the private VPN). The WAN ports' addresses can be allocated by the service providers or dynamically assigned (e.g., by DHCP).



CN: Client Network

Figure4: SD-WAN with Internet facing ports (A1,B2,C3,D2)

Services may not be congruent, i.e., the packets from A-> B may traverse one underlay network, and the packets from B -> A may go over a different underlay due to SD-WAN policies being applied independently in each direction.

3.4. Scenario #3: Private VPN PE based SD-WAN

Private VPN PE-Based SD-WAN refers to extending an existing VPN (e.g., EVPN [RFC7432] or IPVPN) by adding additional ports that face the public Internet to Provider Edge (PE) devices. In this scenario, the PE is part of the SD-WAN edge function, enabling selected traffic to be offloaded over the public Internet to address increased bandwidth requirements between PE devices. This approach allows VPN service providers to augment their networks without immediately committing to building or leasing new infrastructure.

Key Characteristics of Private VPN PE-Based SD-WAN:

- For MPLS-based VPN, traffic between PEs uses MPLS encapsulation within IPsec tunnels egressing the Internet WAN ports, such as MPLS-in-IP or GRE-in-IPsec.
- The BGP RR remains connected to the PEs via the same trusted network as the original VPN, ensuring consistency in routing policies and security.

The main use case for Private VPN PE-Based SD-WAN is Temporary Bandwidth Expansion.




Figure 5: Additional Internet paths added to the VPN

For Ethernet-based client traffic, Private VPN PE based SD-WAN should support VLAN-based service interfaces (EVPN Instances), VLAN bundle service interfaces, or VLAN-Aware bundling service interfaces. EVPN service requirement as described in Section 3.1 of [RFC8388] are applicable to the SD-WAN Ethernet-based Client services. For IP-based client interfaces, L3VPN service requirements are applicable.

4. Provisioning Model

This section defines the provisioning model, i.e., a conceptual framework that scopes the roles (controller/RR and edges), the artifacts (client-service constructs, policies, and IPsec parameters), and their mapping to BGP attributes.

4.1. Client Service Provisioning Model

Provisioning of client services in an SD-WAN network can leverage approaches similar to those used for VRFs in MPLS based VPNs [RFC4364] [RFC4659]. A client VPN can define communication policies by specifying BGP Route Targets for import and export. Alternatively, policy-based filtering using ACLs (Access Control List) can be employed to control which routes are allowed or denied for a given client VPN.

In scenarios where an SD-WAN edge is dedicated to a single client with a single virtual network, all services attached to the client facing interface(s) on the edge node can be grouped into a single VRF. The RR can manage the policies for import/export policies for that VRF.

4.2. Policy Configuration

Policy configuration is a key characteristic of an SD-WAN service, enabling packets to be forwarded over multiple types of underlays based on predefined rules. Policies determine which underlay paths are allowed to carry specific flows, as outlined in Section 8 of [MEF70.2]. A flow is a collection of packets between the same source and destination pair that are subject to the same forwarding and policy decisions at the ingress SD-WAN edge and are identified by the settings of one or more fields in the packet headers. For example, client-service-x can only be mapped to a

MPLS topology, ensuring traffic alignment with business or security requirements.

4.3. IPsec Related Parameters Provisioning

IPsec-related parameter provision in an SD-WAN network involves the negotiation and distribution of cryptographic parameters required to establish IPsec tunnels among them. To streamline the configuration process, SD-WAN edges can retrieve those parameters directly from the SD-WAN controller, reducing manual intervention and ensuring consistency.

In a BGP-controlled SD-WAN, BGP UPDATE messages could be extended to propagate IPsec-related attributes for each SD-WAN edge [SD-WAN-Discovery]. These attributes allow SD-WAN edges to receive compatible IPsec parameters from the RR over a protected BGP session. By leveraging the authenticated and authorized relationship between each SD-WAN edge and the RR, this approach can simplify IPsec tunnel establishment and reduce or eliminate the need for separate per-peer IKEv2 authentication between SD-WAN edges, depending on the deployment model [RFC7296] (see Section 5.1 and Section 7).

This mechanism supports the ZTP requirement outlined in Section 3.1.4 by enabling IPsec tunnels to be provisioned without IKEv2 negotiation.

5. BGP Controlled SD-WAN

5.1. Rationale for Using BGP as Control Plane for SD-WAN

In small SD-WAN networks with a modest number of nodes, a hub-and-spoke overlay model, including designs that use Next Hop Resolution Protocol (NHRP) [RFC2332] for tunnel endpoint resolution, can be manageable. In larger SD-WAN networks with many edge nodes and diverse underlays, additional mechanisms are often needed to distribute reachability, tunnel endpoint information, and policy-related attributes in a scalable and consistent manner.

BGP offers several key advantages when used as the control plane for a large SD-WAN:

- Simplified peer authentication process:

With a secure channel established between each edge node and its RR, the RR can perform peer authentication on behalf of the edge node. The RR has policies on peer communication and the built-in

capability to constrain the propagation of the BGP UPDATE messages to the authorized edge nodes only.

- Scalable IPsec tunnel management

In networks with multiple IPsec tunnels between SD-WAN edges, BGP can simplify tunnel management by advertising WAN-port properties and IPsec-related parameters. One possible approach for carrying such information is described in [SD-WAN-Discovery]. These BGP attributes allow peers to learn tunnel endpoints and associated parameters via the control plane and to associate advertised client routes with the appropriate IPsec SAs, thereby reducing manual configuration and enabling scalable, consistent tunnel establishment across the SD-WAN.

When an IGP such as OSPF is run over IPsec tunnel interfaces, routing adjacencies are typically established per tunnel. In a BGP-controlled SD-WAN context, multiple service flows can be associated with shared tunnel parameters, reducing repeated per-tunnel control-plane configuration.

- Simplified traffic selection configurations

BGP can simplify the configuration of IPsec tunnel associations and related forwarding policies. By leveraging Route Targets to identify SD-WAN VPN membership, administrators can apply import/export policies that control the distribution of client routes. These route attributes, in turn, inform the local configuration of IPsec traffic selectors at each SD-WAN edge.

- Centralized Management and Security

When the BGP RR is integrated with the SD-WAN controller, it supports a centralized model for managing route distribution policies. The RR ensures that BGP UPDATE messages are distributed only to authorized SD-WAN edges based on preconfigured policies, thereby reducing control-plane complexity and limiting exposure compared to decentralized architectures.

In summary, BGP is well suited to these SD-WAN overlay control-plane functions because it provides standardized mechanisms for scalable route distribution, policy-based route propagation, VPN membership control using Route Targets, and extensible encoding of tunnel-related attributes, which are especially useful as the network grows in size and complexity.

5.2. BGP Scenario for Homogeneous Encrypted SD-WAN

In a BGP-controlled Homogeneous Encrypted SD-WAN, an SD-WAN edge (i.e., C-PE) could advertise both its attached client routes and associated IPsec tunnel parameters using BGP UPDATE messages. One possible approach using the Tunnel Encapsulation Attribute is described in [SD-WAN-Discovery].

For example, in the figure below, the BGP UPDATE message from C-PE2 to RR can and can use BGP attributes to convey the parameters needed to associate those routes with the appropriate IPsec tunnel. One possible approach using the Tunnel Encapsulation Attribute is described in [SD-WAN-Discovery].




Figure 6: Homogeneous Encrypted SD-WAN

In scenarios where C-PE2 does not have a policy specifying the authorized peers for specific client routes, the RR takes the responsibility for ensuring that BGP UPDATE for these client routes are propagated only to other authorized SD-WAN edges.

5.3. BGP Scenario for Differential Encrypted SD-WAN

In this scenario, client services may have distinct forwarding requirements based on business or network policies. Some client services can be routed through any WAN ports of the edge node,

while others must be routed through specific WAN ports (such as only MPLS VPN). To address these requirements, the BGP speaker employs two distinct BGP UPDATE messages:

- UPDATE 1: Client Route Advertisement. This UPDATE advertises the prefixes of client services attached to the client-facing interfaces of an SD-WAN edge. As described in Section 8 of [RFC9012], recursive next-hop resolution can be used to associate each client route with the appropriate WAN-facing interface(s) of the advertising SD-WAN edge for the desired underlay path. Figure 6 illustrates the relationship among the client routes, the SD-WAN edge, and the RR.
- UPDATE 2: WAN-Facing Interface Advertisement. This UPDATE advertises information about the WAN-facing interfaces of an SD-WAN edge that connect to the underlay networks, including attributes such as IPsec SA parameters, MPLS label stacks, and other relevant underlay properties. These attributes are carried in the BGP Tunnel Encapsulation Attribute, together with associated Color values, allowing BGP receivers to associate the advertised WAN-facing interfaces with the client routes advertised in UPDATE 1. Figure 6 also illustrates the SD-WAN edge WAN ports that are the basis for this advertisement.

This dual-update approach provides flexibility and efficiency in managing IPsec tunnels terminated at the WAN ports of SD-WAN edges. By decoupling client route advertisements from IPsec tunnel attributes, it accommodates the differing update frequencies between these components—for example, client route changes may occur independently of dynamic IPsec parameters such as key values. Additionally, multiple client services can share a single IPsec SA, optimizing resource usage and reducing control-plane overhead.

BGP receivers associate the two UPDATE messages using the common loopback address of the SD-WAN edge (e.g., C-PE2). UPDATE 1 advertises client routes with the next-hop set to C-PE2's loopback address. UPDATE 2 advertises underlay WAN port information using an NLRI that contains the same loopback address, along with the Tunnel Encapsulation Attribute conveying IPsec parameters and WAN port properties. BGP receivers use the common loopback address to match the next-hop in UPDATE 1 with the NLRI in UPDATE 2. This enables recursive resolution, as specified in [RFC9012], allowing client traffic to be forwarded based on the underlay characteristics defined in UPDATE 2.

5.4. BGP Scenario for Flow-Based Segmentation

In a flow-based segmentation scenario, as described in Section 3.1.3, a service flow is identified by specific fields in the packet's IP header, such as source/destination IP addresses, port numbers, or protocol types. Flow-based segmentation ensures that traffic for a particular service flow is directed only to authorized nodes or paths, meeting security and policy requirements.

This can be achieved by constraining the propagation of BGP UPDATE messages to nodes that meet the criteria of the service flow. For instance, to enforce communication exclusively between the Payment Application in branch locations and the Payment Gateway, as depicted in Figure 6, the following BGP UPDATE messages can be advertised:

BGP UPDATE #1a: Originated by the SD-WAN edge attached to the Payment Application and propagated only to the Payment Gateway node for a point-to-point (P2P) topology between the Payment Application and the Payment Gateway.

BGP UPDATE #1b: Originated by the same SD-WAN edge for other prefixes and propagated to C-PE1 and C-PE3 for other prefixes that can be reached by these edge nodes.




Figure 7: Flow Based SD-WAN Segmentation

In Figure 7, the "Blue Tunnel" denotes the tunnel used exclusively for Payment Application traffic toward the Payment Gateway, whereas the "Red Tunnels" denote tunnels used for other authorized traffic among SD-WAN edge nodes.

6. SD-WAN Forwarding Model

This section briefly illustrates how BGP-distributed reachability, tunnel attributes, Route Targets, and policy information are used by SD-WAN edges when forwarding client traffic in the scenarios described in Section 3. The forwarding behavior described in this section is not specific to BGP-based control and is provided here only to show how the BGP-distributed information is consumed by SD-WAN edges.

The forwarding procedures described in Section 6 of [RFC8388] are applicable for the SD-WAN client traffic. Similar to the BGP-based VPN/EVPN client routes UPDATE message, Route Targets can be used to distinguish routes from different clients.

6.1. Forwarding Model for Homogeneous Encrypted SD-WAN

6.1.1. Network and Service Startup Procedures

In the Homogeneous Encrypted SD-WAN scenario, two IPsec SAs are required to secure bidirectional traffic between two C-PE nodes (or their client-facing interfaces), since each SA protects traffic in only one direction.

For example, in the full mesh scenario in Figure 3 of Section 3.2, where client CN2 is attached to C-PE1, C-PE3, and C-PE4, six unidirectional IPsec SAs must be established: C-PE1 <-> C-PE3; C-PE1 <-> C-PE4; C-PE3 <-> C-PE4.

SD-WAN services to clients can be IP-based or Ethernet-based. For IP-based services, an SD-WAN edge can learn client addresses from the client-facing interfaces via OSPF [RFC2328] [RFC5340], RIP [RFC1058] [RFC2453], BGP[RFC4271], or static configuration. For Layer-2 services, the EVPN parameters, such as the ESI (Ethernet Segment Identifier), EVI (Ethernet Virtual Instance), and CE-VID (Customer Edge Virtual Instance Identifier) to EVI mapping, can be configured as described in [RFC8388].

In a BGP-controlled SD-WAN, the BGP RR can propagate the client route UPDATE messages to authorized SD-WAN edges based on configured policies. The SD-WAN edges use BGP attributes-such as the Tunnel Encapsulation Attribute and associated Color values-to associate received client routes with the appropriate IPsec Security Associations (SAs), thereby reducing per-edge configuration of tunnel endpoint associations. Other SD-WAN service policies, traffic-steering rules, and device-specific parameters may still be provisioned by the SD-WAN controller or local configuration.

6.1.2. Packet Walk-Through

For unicast packets forwarding:

An IPsec SA terminated at a C-PE node can carry traffic for multiple client services. Packets to and from these services are encapsulated in an inner tunnel, such as GRE or VXLAN. Different client traffic can be distinguished using a unique key or identifier in the inner encapsulation header. This inner tunnel is further encapsulated with an outer IP header, where the source and destination addresses are the loopback addresses of the C-PE nodes, and the protocol field is typically set to ESP (50).

C-PE Node-based IPsec tunnel is inherently protected when the C-PE has multiple WAN ports to different underlay paths. As shown in Figure 2, when one of the underlay paths fails, the IPsec tunnel can continue operating by rerouting traffic through an alternate WAN port.

When a C-PE receives an IPsec encrypted packet from its WAN ports, it decrypts the packet and forwards the inner packet to the client facing interface based on the inner packet's destination address.

For multicast packets forwarding:

While IPsec can support multicast (e.g., using group keying mechanisms), many SD-WAN deployments involve a large number of edge nodes with policy-driven communication patterns, making group key management complex. Therefore, a straightforward and commonly used approach is to encapsulate multicast packets in separate unicast IPsec SA tunnels. More optimized multicast forwarding approaches are outside the scope of this document.

6.2. Forwarding Model for Hybrid Underlay SD-WAN

In this scenario, as shown in Figure 4 of Section 3.3, traffic forwarded over the trusted VPN paths can be native (i.e., unencrypted). The traffic forwarded over untrusted networks need to be protected by IPsec SA.

6.2.1. Network and Service Startup Procedures

Infrastructure setup: The proper MPLS infrastructure must be configured among the edge nodes, i.e., the C-PE1/C-PE2/C-PE3/C-PE4 of Figure 3. The IPsec SA between WAN ports or nodes must be set up as well. IPsec SA related attributes on edge nodes can be distributed by BGP UPDATE messages as described in Section 5.

There could be policies governing how flows can be forwarded, as specified by [MEF70.2]. For example, "Private-only" indicates that the flows can only traverse the MPLS VPN underlay paths.

6.2.2. Packet Walk-Through

Unicast packets forwarding:

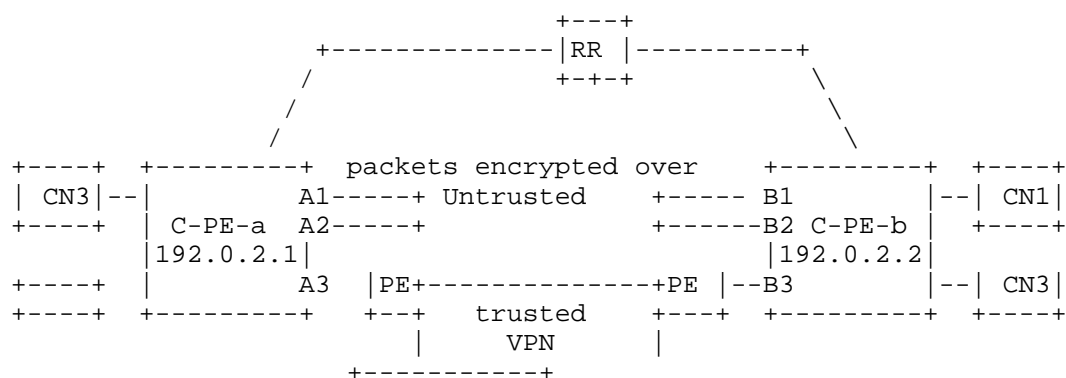
When C-PE-a in Figure 7 receives a packet from a client facing interface, the forwarding decision depends on the flow's routing

policy. If a packet belonging to a flow that must be forwarded over the MPLS VPN, the forwarding processing is the same as the MPLS VPN. Otherwise, C-PE-a can select the least cost path, including the previously established MPLS paths and IPsec Tunnels, to forward the packet. Packets forwarded over the trusted MPLS VPN do not require additional encryption, while those sent over untrusted networks must be encrypted by IPsec SA.

For a C-PE with multiple WAN ports provided by different NSPs, separate IPsec SAs can be established for the WAN ports. In this case, the C-PE have multiple IPsec tunnels in addition to the MPLS path to choose from to forward the packets from the client facing interfaces.

If the IPsec SA is chosen, the packet is encapsulated by the IPsec header and encrypted by the IPsec SA before forwarding it to the WAN.

Packets received over MPLS paths are processed as in standard MPLS VPNs. For packets encrypted with IPsec received from WAN ports, the C-PE decrypts and decapsulates the inner payload before forwarding it according to the local forwarding table. To protect against potential attacks, traffic received through Internet-facing WAN ports should be protected by appropriate security mechanisms (e.g., anti-DDoS, ACLs, or rate-limiting); the specifics of these protections are beyond the scope of this document. Additionally, the control plane must avoid learning routes from Internet-facing WAN ports to ensure network integrity.



Multicast packets forwarding:

For multicast traffic, MPLS multicast [RFC6513], [RFC6514], or [RFC7988] can be utilized to forward multicast traffic across the network.

If IPsec tunnels are used for multicast traffic, the packet must be encapsulated and encrypted separately for each destination, creating multiple unicast IPsec tunnels to deliver the multicast packet to all intended recipients. This approach may have scalability implications due to per-destination packet replication; optimization mechanisms are outside the scope of this document.

6.3. Forwarding Model for PE based SD-WAN

6.3.1. Network and Service Startup Procedures

In this scenario, all PEs have secure interfaces facing the clients and facing the MPLS backbone. Some PEs have additional interfaces to the untrusted public Internet which are for offloading low priority traffic when the MPLS paths get congested. The PEs are already connected to their RRs, and the configurations for the clients and policies are already established.

6.3.2. Packet Walk-Through

When offloading MPLS packets to the Internet path, each MPLS packet is encapsulated by an outer IP header as MPLS-in-IP or MPLS-in-GRE [RFC4023]. The outer IP address can be an interface address or the PE's loopback address.

When IPsec Tunnel mode is used to protect an MPLS-in-IP packet, the entire MPLS-in-IP packet is placed after the IPsec tunnel header. In IPsec transport mode, the MPLS-in-IP packet's IP header becomes the outer IP header of the IPsec packet, followed by an IPsec header, and then followed by the MPLS label stack. The IPsec header must set the payload type to MPLS by using the IP protocol number specified in section 3 of [RFC4023]. For the MPLS-in-GRE packets protected by IPsec Transport Mode, the GRE header follows the IPsec header.

The IPsec SA's endpoints should not be the client-facing interface addresses unless the traffic to/from those clients always goes through the IPsec SA even when the MPLS backbone has enough capacity to transport the traffic.

When the PEs' Internet-facing ports are behind the NAT [RFC3715], additional measures are necessary to support NAT traversal. In this case, an outer UDP field is added to the encrypted payload [RFC3948]. Three specific ports and protocols must remain open on the PEs: UDP port 4500 (used for NAT traversal), UDP port 500 (used for IKE), and IP protocol 50 (ESP). The IPsec IKE (Internet Key Exchange) sessions between PEs navigate NAT environment using the mechanisms outlined in [RFC3947].

When a packet is received from a client facing interface, it is initially processed according to the MPLS VPN forwarding rules. If the MPLS backbone path to the destination is congested, the packet is encapsulated as an MPLS-in-IP packet and encrypted using the IPsec tunnel to the target PE. Conversely, when a packet is received from an Internet-facing WAN port, it is decrypted, and the inner MPLS payload is extracted and forwarded to the MPLS VPN engine for further processing.

As in Scenario #2 (Section 3.3), traffic received from Internet-facing WAN ports should be protected by appropriate security mechanisms. In addition, the control plane should not learn routes from the Internet-facing WAN ports. Operators should take encapsulation overhead and MTU considerations into account when deploying combinations of MPLS-in-IP, GRE, and IPsec encapsulations.

7. Manageability Considerations

A BGP-controlled SD-WAN uses RR to propagate client routes and underlay tunnel properties among authorized SD-WAN edges. Since the RR is configured with policies that identify authorized peers, the peer-wise IPsec IKE (Internet Key Exchange) authentication process is significantly simplified. Operators should ensure that RR propagation policies, SD-WAN VPN membership rules, and tunnel-attribute distribution are consistently configured, as misconfiguration could expose edge properties or cause incorrect binding of client routes to underlay tunnels. While centralizing control in the RR simplifies management, it also means that errors in RR policy or provisioning may have network-wide effects; therefore, careful monitoring, auditing, and validation of RR configuration are essential.

In addition to route and tunnel dissemination, the manageability of a BGP-controlled SD-WAN also encompasses the consistent configuration of segmentation policies, WAN-port attributes, and forwarding behaviors described throughout this document, all of

which rely on the RR as the central point for distributing and validating operational state.

8. Security Considerations

In a BGP-controlled SD-WAN network, secure operation relies in part on the correct configuration and behavior of the RR, which acts as the central distribution point for BGP routing information. The RR applies preconfigured routing policies to control the propagation of BGP UPDATE messages to authorized SD-WAN edges, minimizing the risk of unintended route exposure or unauthorized communication.

SD-WAN operation relies on the existing security mechanisms defined for BGP and IPsec, which therefore apply directly to the control and tunnel planes described in this document. A primary operational difference between SD-WAN deployments and BGP/MPLS IP VPNs [RFC4364][RFC4659] is that SD-WAN edge nodes often include Internet-facing WAN ports, which introduce additional security, filtering, and policy-enforcement requirements that are not typically present in BGP/MPLS IP VPN environments. These untrusted interfaces increase exposure to spoofed traffic, denial-of-service attacks, and unintended route learning if misconfigured. As a result, operators must apply strict validation of control-plane information received from Internet-facing ports, ensure correct RR policy configuration, and provide appropriate protection for both control-plane and data-plane exchanges, consistent with the security guidance in the referenced RFCs.

The security model for the SD-WAN described in this document is based on the following principles:

- 1) Centralized Control: The RR governs all routing and policy decisions. This centralized architecture simplifies security management compared to distributed models, as it limits the potential attack surface to a smaller, more controlled set of components. However, this centralization also means that misconfiguration of RR policies or controller logic can have network-wide impact, potentially exposing edge properties to unauthorized peers, allowing incorrect route propagation, or disrupting IPsec tunnel mappings across the entire SD-WAN network.
- 2) Secure Channels: All communication between SD-WAN edges and the RR must occur over a secure channel, such as TLS or IPsec, to ensure the confidentiality and integrity of BGP UPDATE messages.

If a secure communication channel is not used, an attacker on the path could observe or tamper with BGP UPDATES, potentially exposing edge properties or injecting unauthorized routes. Deployments are therefore expected to enforce secure transport (e.g., reject or disable sessions that are not established over a protected channel).

- 3) Policy Enforcement: The RR is responsible for enforcing policies that restrict the propagation of edge node properties and routing updates to only authorized peers. This prevents sensitive information from being exposed to unauthorized nodes. Misconfigured RR policies could inadvertently expose edge properties or client routes to unauthorized peers, or conversely block required updates, leading to traffic disruption or unintended connectivity; these risks highlight the importance of correct and audited policy configuration.
- 4) Mitigation of Internet-facing Risks: In scenarios where SD-WAN edges include Internet-facing WAN ports, additional measures must be taken to mitigate security risks:
 - Anti-DDoS mechanisms must be enabled to protect against potential attacks on Internet-facing ports. For example, rate limiting, basic anomaly detection, or upstream filtering can reduce exposure to volumetric attacks; without such protections, Internet-facing WAN ports may be vulnerable to flooding that can disrupt both data and control plane reachability.
 - The control plane must avoid learning routes from Internet facing WAN ports to prevent unauthorized traffic from being injected into the SD-WAN. If this control-plane filtering is misconfigured, an SD-WAN edge may inadvertently learn and propagate untrusted Internet routes, enabling route injection attacks or unauthorized reachability into the SD-WAN overlay.

By concentrating the security within the RR and using secure channels, the SD-WAN network achieves consistent enforcement of security policies and reduces the likelihood of misconfigurations at individual edge nodes. However, the robustness of this security model depends critically on the proper configuration and ongoing maintenance of the RR. Operators must ensure that the RR itself is

adequately protected against compromise or misconfiguration, as its failure or exploitation could impact the entire network.

This model emphasizes simplicity and efficiency, leveraging centralized governance to mitigate risks while supporting scalable control-plane distribution and interoperability of the SD-WAN.

9. IANA Considerations

No Action is needed.

10. References

10.1. Normative References

- [RFC4301] S. Kent and K. Seo, "Security Architecture for the Internet Protocol", RFC4301, Dec. 2005.
- [RFC4360] S. Sangli, et al, "BGP Extended Communities Attribute", RFC4360, Feb. 2006.
- [RFC4364] E. Rosen, Y. Rekhter, "BGP/MPLS IP Virtual Private networks (VPNs)", Feb 2006.
- [RFC4456] T. Bates, E. Chen, R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", April 2006.
- [RFC4659] J. De clercq, et al, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC4659, Sept 2006.
- [RFC9012] K. Patel, et al "The BGP Tunnel Encapsulation Attribute", RFC9012, April 2021.
- [MPLIFY-119] SD-WAN Service Attributes and Service Framework, <https://www.mplify.net/resources/mplify-119-universal-sd-wan-edge-implementation-agreement/>. June 2025.

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

10.2. Informative References

- [RFC1058] C. Hedrick, "Routing Information Protocol", RFC1058, June 1988.
- [RFC2328] J. Moy, "OSPF Version 2", RFC2328, April 1998.
- [RFC2332] J. Luciani, et al, "NBMA Next Hop Resolution Protocol (NHRP)", RFC2332, April 1998.
- [RFC2453] G. Malkin, "RIP Version 2", RFC2453, Nov. 1998.
- [RFC2784] D. Farinacci, et al, "Generic Routing Encapsulation (GRE)" RFC2784, March 2000.
- [RFC3715] B. Aboba, W. Dixon, "IPsec-Network Address Translation (NAT) Compatibility Requirements", March 2004.
- [RFC3947] T. Kivinen, et al, "Negotiation of NAT Traversal in the IKE", Jan. 2005.
- [RFC3948] A. Huttunen, et al, "UDP Encapsulation of IPsec ESP Packets", Jan 2005.
- [RFC4023] T. Worster, Y. Rekhter, E. Rosen, "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", March 2005.
- [RFC4271] Y. Rekhter, et al, "A Border Gateway Protocol 4 (BGP-4)", RFC4271, Jan 2006.
- [RFC4761] K. Kompella and Y. Rekhter, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC4761, Jan. 2007.
- [RFC4762] M. Lasserre and V. Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC4762, Jan. 2007.

- [RFC5340] R. Coltun, et al, "OSPF for IPv6", RFC5340, July 2008.
- [RFC6513] E. Rosen and R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", RFC6513, Feb, 2012.
- [RFC6514] R. Aggarwal, et al, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC6514, Feb, 2012.
- [RFC7296] C. Kaufman, et al, "Internet Key Exchange Protocol Version 2 (IKEv2)", Oct 2014.
- [RFC7348] M. Mahalingam, et al, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC7348, Aug 2014.
- [RFC7432] A. Sajassi, et al, "BGP MPLS-Based Ethernet VPN", RFC7432, Feb 2015.
- [RFC7988] E. Rosen, et al, "Ingress Replication Tunnels in Multicast VPN", RFC7988, Oct. 2016.
- [RFC8200] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification". July 2017.
- [RFC8388] J. Rabadan, et al, "Usage and Applicability of BGP MPLS-Based Ethernet VPN", RFC8388, May 2018.
- [RFC9522] A. Farrel, "Overview and Principles of Internet Traffic Engineering", RFC9522, Jan. 2024.
- [SD-WAN-Discovery] L. Dunbar, et al "BGP UPDATE for SD-WAN Edge Discovery", draft-ietf-idr-sdwan-edge-discovery-27, April 2026.
- [Net2Cloud-Problem] L. Dunbar and A. Malis, "Dynamic Networks to Hybrid Cloud DCs: Problems and Mitigation Practices", draft-ietf-rtgwg-net2cloud-problem-statement-43, April. 2025.
- [IEEE802.3] "IEEE Standard for Ethernet" by The Institute of Electrical and Electronics Engineers (IEEE) 802.3.

11. Acknowledgments

Acknowledgements to Gunter van de Velde, Andrew Alston, Adrian Farrel, Jim Guichard, Joel Halpern, John Scudder, Darren Dukes, Andy Malis, Donald Eastlake, Stephen Farrell, and Cheng Sheng for their review and contributions.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Linda Dunbar
Futurewei
Email: ldunbar@futurewei.com

Ali Sajassi
Cisco
Email: sajassi@cisco.com

John Drake
Independent
Email: je_drake@yahoo.com

Basil Najem
Bell Canada
Email: basil.najem@bell.ca

Sue Hares
Email: shares@ndzh.com

Contributor's Addresses

David Carrel
Graphiant
Email: carrel@graphiant.com

Ayan Banerjee
Cisco
Email: ayabaner@cisco.com

