

Routing Area Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 15 November 2025

C. Sheng  
H. Shi  
Huawei Technologies  
L. Dunbar  
Futurewei  
J. Dong  
Huawei Technologies  
14 May 2025

Scenarios and Challenges of Overlay Routing for SD-WAN  
draft-sheng-rtgwg-overlay-routing-requirement-02

Abstract

Overlay routing is essential during the enterprise networks' evolution from the interconnection among multiple on-premise branch sites to more advanced ones, such as the interconnection to multi-clouds. This document analyzes the technical requirements and challenges of overlay routing for SD-WAN in these scenarios.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Routing Area Working Group Working Group mailing list ([rtgwg@ietf.org](mailto:rtgwg@ietf.org)), which is archived at <https://mailarchive.ietf.org/arch/browse/rtgwg/>.

Source for this draft and an issue tracker can be found at <https://github.com/VMatrix1900/draft-rtgwg-overlay-routing-requirement>.

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 15 November 2025.

## 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 . . . . .	2
2. Terminology . . . . .	3
2.1. Requirements Language . . . . .	3
3. Virtual tunnel auto-discovery and provision requirement . . .	3
4. Topology-aware routing with multi-hop overlay network requirement . . . . .	5
5. SLA-aware routing across multiple overlay segments requirement . . . . .	5
6. Seamless integration with virtual networks of multiple clouds requirement . . . . .	6
7. Overlay multicast over multicast-free underlay networks requirement . . . . .	7
8. Security Considerations . . . . .	7
9. Acknowledgement . . . . .	7
10. References . . . . .	7
10.1. Normative References . . . . .	7
10.2. Informative References . . . . .	7
Authors' Addresses . . . . .	8

## 1. Introduction

SD-WAN is currently widely used in the basic scenarios of one-hop interconnection between enterprise on-premises sites of branches, campuses, and even DCs. With multi-cloud adoption, workloads are migrating to be hosted on clouds. It is necessary for SD-WAN to interconnect multiple on-premises sites and multiple cloud sites seamlessly with the overlay routing technology.

As the core network technology, overlay routing faces a series of new challenges during its evolution, such as flexible overlay topology formation and auto-provision, global interconnection among multi-regions via multi-ISP networks, and the SLA aware routing across multiple overlay segments. Also, it is necessary to investigate how SD-WAN can be seamlessly integrated with the virtual network of multiple clouds.

## 2. Terminology

- \* SD-WAN: Software Defined Wide Area Network. In this document, "SD-WAN" refers to the policy-driven transporting of IP packets over multiple different underlay networks to get better WAN bandwidth management, visibility, and control.
- \* Site: Enterprise sites across different geo regions, where people or workload host, such as branches, campuses, or even clouds.
- \* Edge: The border devices of the SD-WAN site, which could be physical or virtual CPEs.
- \* TN: Transport Network, the underlay connectivity network which correspond to different ISP network between SD-WAN sites.
- \* TNP: Transport Network Port, the wan port of the Edge which connects to TN.
- \* Virtual Tunnel: The IP tunnel between two TNPs of two different edges from different sites.

### 2.1. Requirements Language

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 BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 3. Virtual tunnel auto-discovery and provision requirement

As the basis of the SD-WAN overlay network, virtual tunnels between edges should be established before the client routes exchange packets. The virtual tunnels, such as IPSec tunnels, establishment between edges require extensive information exchange, such as public keys, tunnel endpoints properties, etc., which can significantly delay client route packet forwarding if they are not established ahead of time. A virtual tunnel is a point-to-point forwarding relationship between two SD-WAN Edges across a given or multiple

underlay ISP networks that provide a well-defined set of transport characteristics (e.g., delay, security, bandwidth, etc.).




Figure 1: Hub spoke topology




Figure 2: Full mesh topology




Figure 3: Layered topology

Different enterprises often have different connectivity topologies with hundreds and thousands of tunnels, as shown in Figure 1, Figure 2 and Figure 3. For the efficiency and simplicity of the O&M, it is highly expected to discover and establish the virtual tunnels between sites automatically instead of manually configuring the overlay tunnels one by one.

[I-D.ietf-idr-sdwan-edge-discovery] has designed an efficient mechanism to exchange the information of each endpoint of the overlay tunnel by BGP protocol, by which edges could check and decide to establish the tunnel or not automatically. While this mechanism works fine in reality, it can be further improved. For example, it is much more expected to carry more information to reflect the topology intent (Full Mesh, P2MP, P2P) in BGP.

#### 4. Topology-aware routing with multi-hop overlay network requirement

There are many differences in the control plane between the traditional L3 VPN network and the SD-WAN overlay network. As per L3VPN network, IGP protocol (OSPF or ISIS) is deployed on each physical link between routers and is responsible for discovering the underlay network topology and calculating the routing of the BGP nexthops (often loopback0 of PEs), while BGP is deployed to advertise and calculate the VPN routes based on the IGP output. In the SD-WAN overlay network, it is difficult and a not good choice to run IGP directly on the tunnels between edges because it will bring much more resource consumption. p2p tunnels, such as GRE or VXLAN, need to be configured using a virtual interface to run the IGP protocol. Flooding of the IGP message could cause resource waste in the control plane.

For the SD-WAN overlay network, it is recommended to use BGP to discover the overlay topology and calculate the best overlay path, which is also responsible for advertising and calculating the VPN routes.

#### 5. SLA-aware routing across multiple overlay segments requirement

After a multi-hop SD-WAN overlay network is established, such as the one shown in Figure 4 below, stitching together the overlay tunnels across the Edge1-Edge2-Edge5-Edge6 for the client traffic between Edge1 and Edge6 might provide better SLA than building other overlay tunnels between Edge1 and Edge6, such as Edge1-Edge2-Edge4-Edge6, etc. Importing traffic engineering based routing in overlay network can provide more deterministic end-to-end QoS SLA for application.




Figure 4: Example of SLA aware routing

Different application flows have different SLA requirements. For example, voice is sensitive to latency and jitter, while video requires a low packet loss forwarding path. It is necessary to provide some degree of TE function to meet the requirements of different types of applications, which is a new challenge for the SD-WAN overlay networks. Naturally, the centralized SD-WAN controller MUST collect SLA (latency, packet loss, and bandwidth) information of the tunnels and the overall topology to calculate the segment list satisfying the requirement raised by the application. Further, the data plane that can carry the overlay tunnel list needs to be carefully designed with the consideration of efficiency and productivity.

#### 6. Seamless integration with virtual networks of multiple clouds requirement

As more and more enterprises migrate their workloads to multiple clouds, it is highly expected to establish a high-quality interconnection between the enterprise's on-premise sites and the cloud sites with qualified O&M specification.

It has been widely adopted to create vCPEs on the clouds as cloud edge to bring a uniform experience and O&M method for access to the clouds. There are also obstacles discovered. For example, how to integrate the multi VPN or multi-tenants to the virtual network of different clouds. And for the sake of reliability, at least two vCPEs need to be created for each cloud site. And it is often recommended to deploy VRRP between the two vCPEs, which need to run the VRRP control plane over multicast packets. For the reason of security, many cloud service providers closed the native IP multicast services for the tenants. So some new HA features need to be considered in such scenarios.

Also, different cloud service provider implements different charge rule for the resources of the compute, network, etc. It needs to be finely scrutinized to develop the most economical network solution for SD-WAN in cloud networks.

## 7. Overlay multicast over multicast-free underlay networks requirement

As more and more enterprise applications are running across SD-WAN overlay networks, multicast traffic is also emerging. Different from traditional multicast VPN networks, SD-WAN overlay networks are based on multiple underlay ISP networks, such as internet, 5G, MPLS, etc which do not support multicast. How to implement a multicast overlay network on top of the multicast-free underlay is challenging. Enhancement to the existing SD-WAN routing protocol needs to be made.

## 8. Security Considerations

TBD

## 9. Acknowledgement

The authors would like to thank Haibo Wang, Shunwan Zhuang, Donglei Pang, Hongwei He for their help.

## 10. References

### 10.1. Normative References

- [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>>.
- [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>>.

### 10.2. Informative References

- [I-D.ietf-idr-sdwan-edge-discovery] Dunbar, L., Hares, S., Majumdar, K., Raszuk, R., and V. Kasiviswanathan, "BGP UPDATE for SD-WAN Edge Discovery", Work in Progress, Internet-Draft, draft-ietf-idr-sdwan-edge-discovery-22, 17 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sdwan-edge-discovery-22>>.

Authors' Addresses

Cheng Sheng  
Huawei Technologies  
China  
Email: shengcheng@huawei.com

Hang Shi  
Huawei Technologies  
China  
Email: shihang9@huawei.com

Linda Dunbar  
Futurewei  
Email: linda.dunbar@futurewei.com

Jie Dong  
Huawei Technologies  
China  
Email: jie.dong@huawei.com