Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: 18 November 2025

C. Sheng
H. Shi
J. Dong, Ed.
Huawei
L. Dunbar
Futurewei
G. Mishra
Verizon

17 May 2025

Associated Gateway Exchange in Multi-segment SD-WAN
draft-sheng-idr-gw-exchange-in-sd-wan-03

Abstract

The document describes the control plane enhancement for multi-segment SD-WAN to exchange the associated GW information between edges.

Discussion Venues

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

Discussion of this document takes place on the Inter-Domain Routing Working Group mailing list (idr@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/idr/.

Source for this draft and an issue tracker can be found at https://github.com/VMatrix1900/draft-sheng-idr-gw-exchange-in-sd-wan.

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 18 November 2025.

Table of Contents

1. Introduction

[I-D.draft-dmk-rtgwg-multisegment-sdwan] describes how enterprises leverage Cloud Providers' backbone infrastructure to interconnect their branch offices. As illustrated in Figure 1, CPE-1 and CPE-2 establish connections to their respective Cloud Gateways (GW) in distinct regions. CPE-1 and CPE-2 maintain the pairwise IPsec Security Associations (SAs). The IPsec encrypted traffic from CPE-1 to CPE-2 is encapsulated by the GENEVE header [RFC8926], with the outer destination address being the GW1.

[I-D.draft-dmk-rtgwg-multisegment-sdwan] specifies a set of sub-TLVs to convey information about the GWs associated with the destination branches, such as GW3 for CPE-2, along with additional attributes. To accomplish this, CPE-1 must be aware of the associated GW addresses of their peers. This document proposes a BGP extension building upon [I-D.draft-ietf-idr-sdwan-edge-discovery], enabling a CPE to advertise its directly connected GW address to other CPEs.

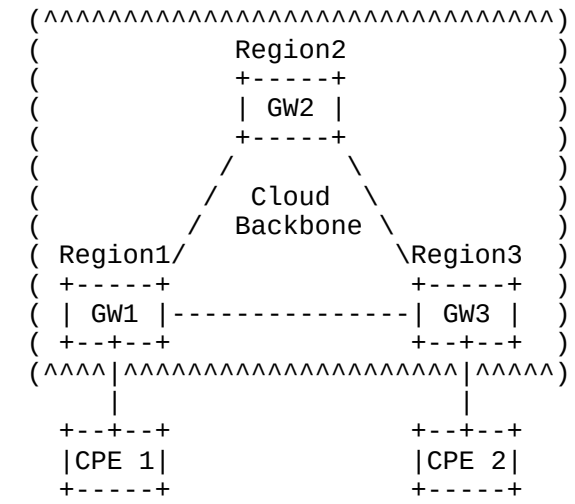

Figure 1: Multi-segment SD-WAN

2. 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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The following acronyms and terms are used in this document:

* Cloud DC: Off-Premises Data Center, managed by 3rd party, that hosts applications, services, and workload for different organizations or tenants.

* CPE: Customer (Edge) Premises Equipment.

* OnPrem: On Premises data centers and branch offices.

* RR: Route Reflector.

* SD-WAN: Software Defined Wide Area Network. In this document, “SD-WAN” refers to a policy-driven transporting of IP packets over multiple underlay networks for better WAN bandwidth management, visibility, and control.

* VPN: Virtual Private Network.

3. Extension to SD-WAN Underlay UPDATE for Associated GWs

The Client Routes Update is the same as described in Section 5 of [I-D.draft-ietf-idr-sdwan-edge-discovery].

3.1. NLRI SD-WAN SAFI Route Type For GW

Adding a new attribute (Associated-Gateway Sub-TLV) to the SD-WAN-Hybrid Tunnel Encoding which is included in the SD-WAN SAFI (=74) Underlay Tunnel Update:

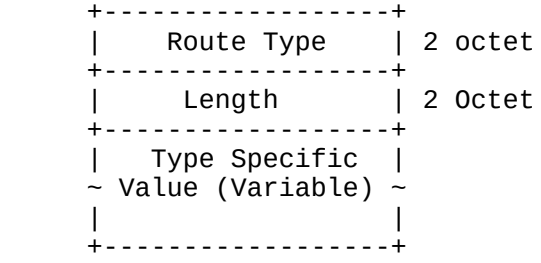

NLRI Route-Type = 2: For advertising the detailed properties of the transit gateways for the edge. The SD-WAN NLRI Route-Type =2 has the following encoding:

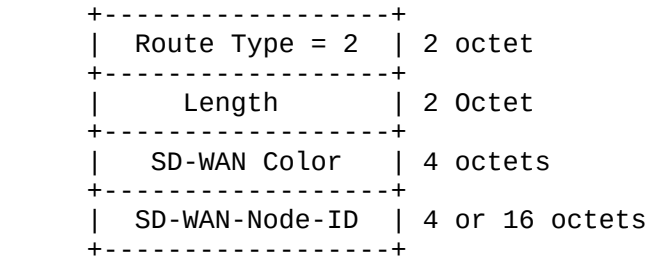

SD-WAN-Color: To represent a group of tunnels that correlate with the Color-Extended-community included in a client route UPDATE. When multiple SD-WAN edges can reach a client route co-located at one site, the SD-WAN- Color can represent a group of tunnels terminated at those SD-WAN edges co-located at the site, which effectively represents the site.

SD-WAN Node ID: The node’s IPv4 or IPv6 address.

3.2. Associated GW Sub-TLV

The Associated GW Sub-TLV, within the SD-WAN-Hybrid Tunnel TLV (code point 25), carries the associated GW address(es) with which the SD-WAN edge is associated.

The following is the structure of the associated GW Sub-TLV:

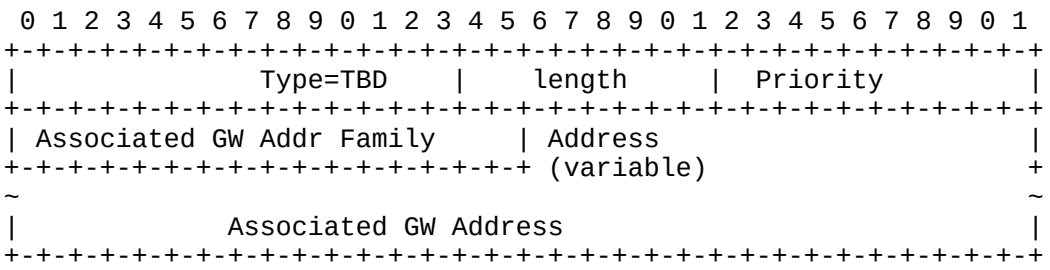

Priority (1-255) indicate the preference of the GW. The higher the value, the more preference is the GW.

Priority = 0 represents that the connection exists but request Cloud Backbone not to choose the GW if possible.

4. Manageability Considerations

Effective management of SD-WAN edge nodes and the exchange of associated cloud gateway information are critical aspects in ensuring a robust and scalable SD-WAN deployment.

5. Security Considerations

This document does not introduce any new security considerations.

6. IANA Considerations

Need IANA to assign a new Sub-TLV Type under the SD-WAN-Hybrid Tunnel TLV.

* SD-WAN Associated GW Sub-TLV.

7. Normative References

[I-D.draft-dmk-rtgwg-multisegment-sdwan]
Majumdar, K., Dunbar, L., Kasiviswanathan, V., Ramchandra, A., and A. Choudhary, "Multi-segment SD-WAN via Cloud DCs", Work in Progress, Internet-Draft, draft-dmk-rtgwg-multisegment-sdwan-07, 20 February 2024, <https://datatracker.ietf.org/doc/html/draft-dmk-rtgwg-multisegment-sdwan-07>.

[RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <https://www.rfc-editor.org/rfc/rfc8926>.

[I-D.draft-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>.

[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/rfc/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/rfc/rfc8174>.

Authors' Addresses

Cheng Sheng
Huawei
Beiqing Road
Beijing
Email: shengcheng@huawei.com

Hang Shi
Huawei
China
Email: shihang9@huawei.com

Jie Dong (editor)
Huawei
China

Email: jie.dong@huawei.com

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

Gyan Mishra
Verizon
Email: gyan.s.mishra@verizon.com