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

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

17 May 2025

Advertising SaaS Path Performance Metrics using BGP
draft-sheng-idr-advertising-saas-path-performance-01

Abstract

This document extends BGP to advertise the SaaS path performance metrics from the gateway sites to branch sites. The user can access SaaS applications through the DIA (Direct Internet Access) link at the branch site or through the DIA link at the gateway site, or use the DIA link of a gateway site for redundancy. This approach will improve the SaaS access experience for end-users.

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-advertising-saas-path-performance.

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

With the continuous cloudification of enterprise IT architectures and widespread use of public clouds, more and more enterprises are turning their infrastructures (such as enterprise data centers) to cloudification, abandoning traditional closed IT architectures and using open network architectures. To further achieve this goal, enterprises' mission-critical applications, such as office, production ERP systems, and sales systems, are migrated to the cloud.

In this case, enterprises increasingly rely on software as a service (SaaS) provided by application service providers and prefer to access mission-critical applications from the cloud over the Internet.

Accessing SaaS applications like SalesForce, SharePoint, Dropbox and Office 365 over congested public networks can be unreliable and slow, due to heavy traffic, packet loss, and fluctuating latencies. Application slowness results in poor end-user experience.

This document provides a way to improve the SaaS access experience. As shown in the Figure 1, user can access SaaS applications through the DIA (Direct Internet Access) link at the branch site or through the DIA link at the gateway site. The GWs at the gateway site normally have stronger capabilities and will provide SaaS access services for branch sites. The CPE at the branch site need to choose the best path for each SaaS application. The performance of the path between gateway and SaaS application needs to be advertised to CPE. This document extends BGP to advertise the SaaS path performance metrics.

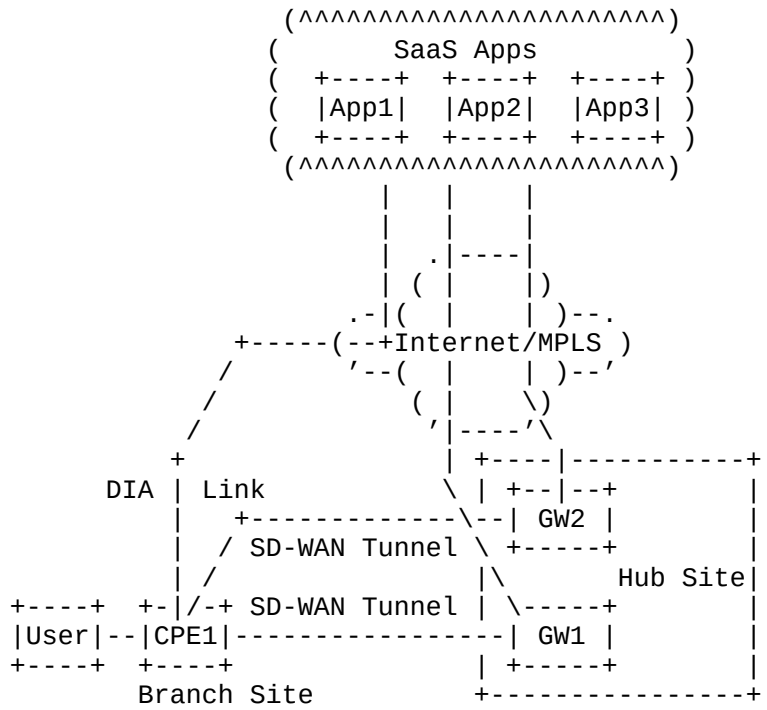

Figure 1: SaaS Application Path Performance Optimization Scenario

2. Terminology

In addition to terms defined in [I-D.ietf-idr-sdwan-edge-discovery], this document uses following terms:

* DIA: Direct Internet Access
* FQDN: Fully Qualified Domain Name
* QoS: Quality of Service
* SaaS: Software-as-a-Service

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

3. Dynamically Select the Best Path

This section uses the scenario shown in Figure 1 as an example to describe how to implement the SaaS Path Optimization solution.

Both the Branch and GW routers initiate periodic probes to target SaaS applications. The GW routers advertise the probe result to the Branch routers. The following figure shows the SaaS Path Performance Metrics table on the CPE1. Note that in this example, CPE1, GW1, and GW2 have multiple paths for accessing App1, 2 paths are listed for each device. The access to App2 and App3 is similar, only one entry is listed for the purpose of simplifying the description.

Name	ID	Path Index	Path Out Intf (# Remote)	O_QoS	Status	L	D	J	B	F_QoS
App1	10	I11	GE 0/0/1.1	75	Good	1	150	40	B01	75
App1	10	I12	GE 0/0/1.1	80	Good	1	160	40	B01	80
App1	10	I13	# GW1-System IP	85	Good	0	100	40	B11	83
App1	10	I14	# GW1-System IP	85	Good	0	100	40	B12	81
App1	10	I15	# GW2-System IP	90	Best	0	80	20	B13	82
App1	10	I16	# GW2-System IP	90	Best	0	80	20	B14	88
App2	20	I02	GE 0/0/1.1	40	Issue	5	180	101	B02	40
App2	20	I21	# GW1-System IP	80	Good	1	100	70	B21	75
App2	20	I22	# GW2-System IP	60	Acct	3	160	80	B22	55
App3	30	I03	GE 0/0/1.1	90	Best	0	58	20	B03	90
App3	30	I31	# GW1-System IP	80	Good	0	65	30	B31	78
App3	30	I32	# GW2-System IP	75	Acct	2	130	90	B32	72

L: Loss D: Delay
J: Jitter B: Bandwidth
Acct: Acceptable O_QoS: Original QoS
F_QoS: Final QoS

Figure 2: CPE1’s SaaS Path Performance Metrics Table

Upon receiving the QoS score from the GW router, CPE1 will calculates the Final QoS score based on the SD-WAN tunnel status and and the received QoS score. When a user of CPE1 accesses a SaaS applications, CPE1 determines the best performing path toward the SaaS application based on the Final QoS score (F_QoS).

For example If App1 is the target SaaS Application, select the SaaS path that passes through GW2 with the Path Index I16 because it has the highest score: 88. If App2 is the target SaaS Application, select the SaaS path that passes through GW1 with the Path Index I21 because it has the highest score: 75. If App3 is the target SaaS application, select the local SaaS path with the Path Index I03 because it has the highest score: 90.

4. The SaaS Path Performance Route

The BGP SD-WAN NLRI as defined in [I-D.ietf-idr-sdwan-edge-discovery] is shown below:

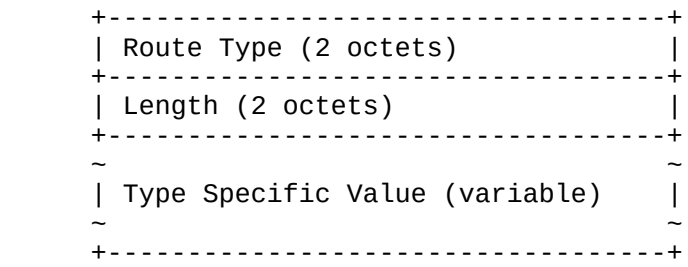

Figure 3: BGP SD-WAN NLRI

Where:

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

* Length: 2 octets of length expressed in bits as defined in [RFC4760].

This document defines an additional route type to be used for the advertisement of the SaaS Path Performance Metrics between different enterprise sites:

* NLRI Route Type: 2

* Name: SaaS Path Performance Route

4.1. The SaaS Path Performance Route Encoding

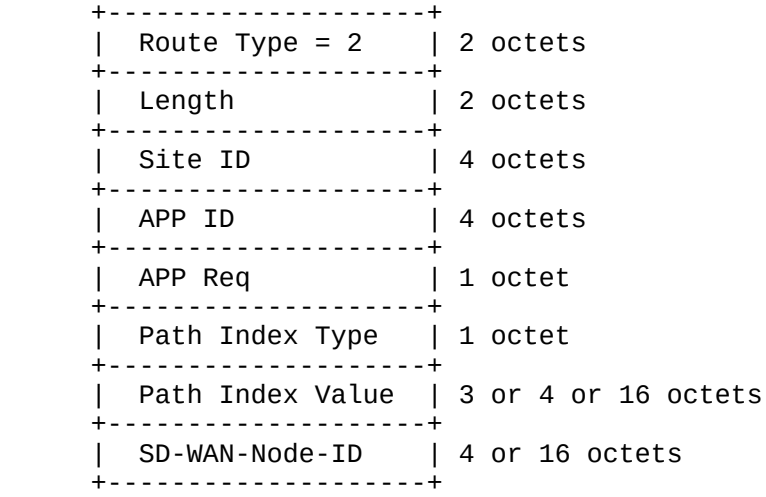

Figure 4: SaaS Path Performance Route

Where:

* Route Type: 2, SaaS Path Performance Route

* Length: 2 octets of length expressed in bits as defined in [RFC4760].

* Site ID: 4 octets, A site ID is a unique identifier of an enterprise site in the SD-WAN network.

* APP ID: 4 octets, SaaS Application ID, a unique Application ID to identify different applications. Application may be deployed using different IP address in different area. Thus an ID is needed to identify the application.

* APP Req: 1 octet, Application requirement to indicate the application requirement of the path quality. For example, an real time video conferencing application requires higher quality than a background file backup application. The value includes:

  - Type = 1: default;

  - Type = 2: Medium;

  - Type = 3: High;

* Path Index Type: Indicates the type of the path index.

* Path Index Value: a Path Index Type specific Value:

  - Type 1, the Path Index Value is a 4-byte local index value, which is used to identify an outbound interface for accessing SaaS applications.

  - Type 2, the Path Index Value is a 3-byte MPLS label, which is used to identify an outbound interface for accessing the SaaS application.

  - Type 3, The Path Index Value is a 16-byte SRv6 SID, which is used to identify an outbound interface for accessing a SaaS application, and its Endpoint Behavior is End.DT2SaaSPath: Decapsulate SRv6 packet, then send the packet to the target SaaS application from the outbound interface indicated by the SRv6 SID.

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

4.2. The SaaS Path Performance Metrics Encoding

The Metadata Path Attribute has been as defined in [I-D.ietf-idr-5g-edge-service-metadata]. This document introduces some additional Sub-TLVs to encode the SaaS Path Performance Metrics and SaaS Application Information.

Another option is to use the above Sub-TLVs in the Tunnel Encapsulation Attribute [RFC9012]. In this option, the tunnel type "SaaS Application Path Performance" is added.

4.2.1. The SaaS Path Delay Sub-TLV format

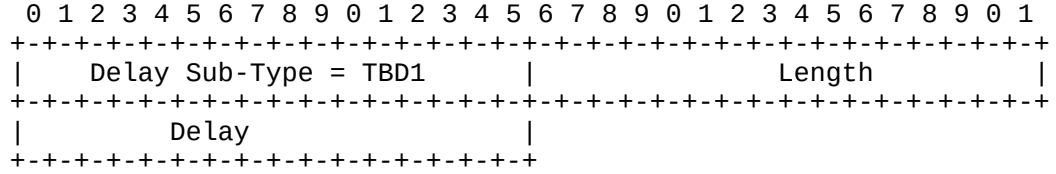

Figure 5: SaaS Path Delay Sub-TLV

Where:

* Delay Sub-Type: TBD by IANA.

* Length: 2 octets, the total number of octets of the value field.

* Delay: 2 octets, this field indicates the packet transmission delay, in milliseconds.

4.2.2. The SaaS Path Loss Sub-TLV format

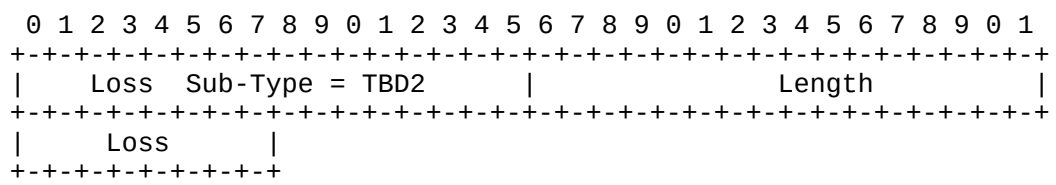

Figure 6: SaaS Path Loss Sub-TLV

Where:

* Loss Sub-Type: TBD by IANA

* Length: 2 octets, the total number of octets of the value field.

* Loss: 1 octet, this field indicates the packet loss rate (%).

4.2.3. The SaaS Path Jitter Sub-TLV format

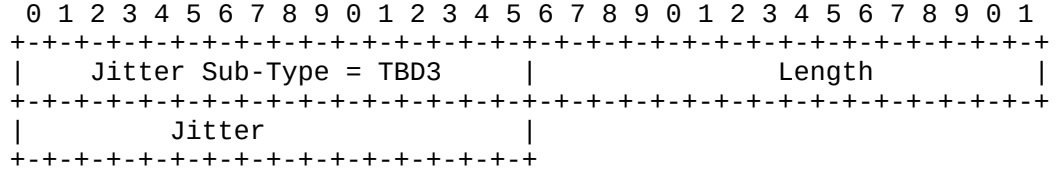

Figure 7: SaaS Path Jitter Sub-TLV

Where:

* Jitter Sub-Type: TBD by IANA

* Length: 2 octets, the total number of octets of the value field.

* Jitter: 2 octets, this field indicates the jitter on the SaaS Path. Range: 1 through 1000 milliseconds

4.2.4. The SaaS Path Bandwidth Sub-TLV format

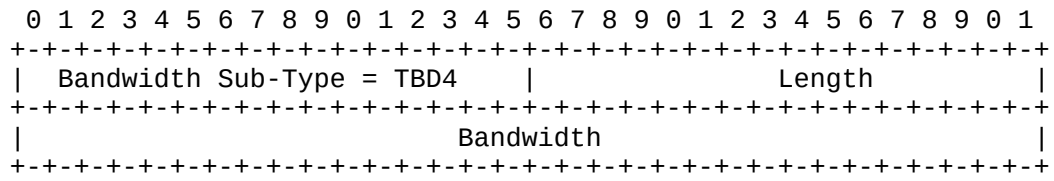

Figure 8: SaaS Path Bandwidth Sub-TLV

Where:

* Bandwidth Sub-Type: TBD by IANA

* Length: 2 octets, the total number of octets of the value field.

* Bandwidth: 4 octets, this field indicates the bandwidth of the SaaS Path.

4.2.5. The SaaS Path Status Sub-TLV format

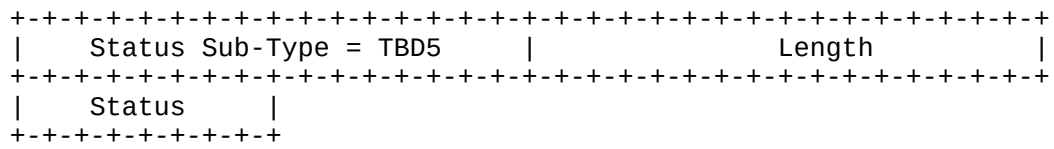

Figure 9: SaaS Path Status Sub-TLV

Where:

* Status Sub-Type: TBD by IANA

* Length: 2 octets, the total number of octets of the value field.

* Status: 1 octet, Network assessment, there are 6 levels as follows:

  - 100: Best

  - 80: Good, Meets recommendations

  - 60: Acceptable

  - 40: Users may experience issues

  - 20: Users may complain

- 0: Network problems

4.2.6. The SaaS Path QoS Sub-TLV format

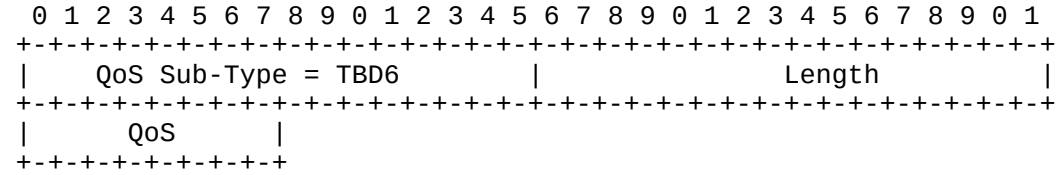

Figure 10: SaaS Path QoS Sub-TLV

Where:

* QoS Sub-Type: TBD by IANA

* Length: 2 octets, the total number of octets of the value field.

* QoS: 1 octet, Quality of Service, 1-100, with 1 being the worst, and 100 being the best. The QoS value is calculated based on the values of Loss, Jitter, Delay, and Status.

4.2.7. The SaaS Application Name Sub-TLV format

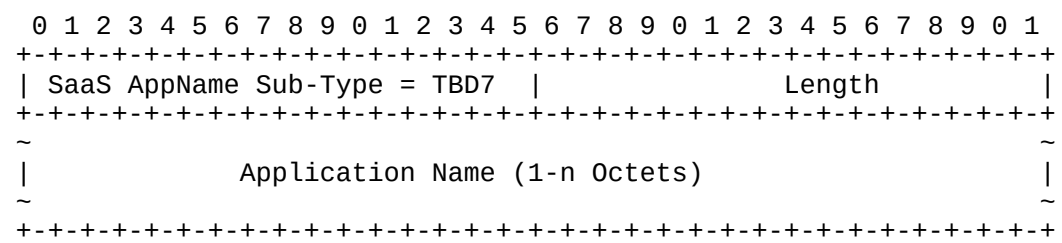

Figure 11: Saas Application Name Sub-TLV

Where:

* SaaS AppName Sub-Type: TBD by IANA

* Length: 2 octets, the total number of octets of the value field.

* Application Name: The name of the application represented as a string, such as Salesforce, Dropbox, Office 365, and so on.

4.2.8. The SaaS Application Domain Name Sub-TLV format

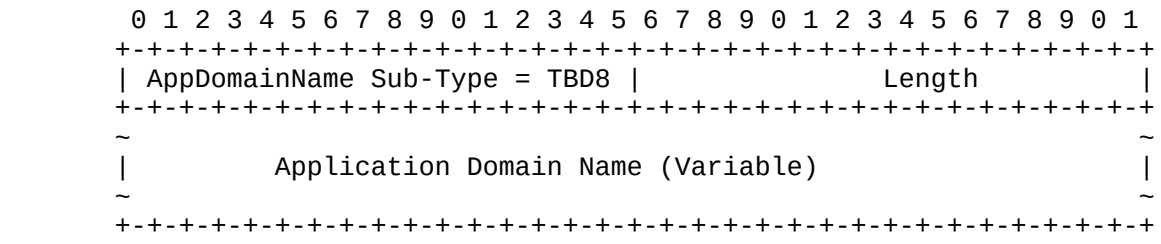

Figure 12: SaaS Application Domain Name Sub-TLV

Where:

* AppDomainName Sub-Type: TBD by IANA

* Length: 2 octets, the total number of octets of the value field.

* Application Domain Name: The domain name of the application represented as a string, such as www.salesforce.com, www.baidu.com, www.iana.org, www.dropbox.com, www.microsoft.com, and so on.

5. Security Considerations

TBD.

6. IANA Considerations

TBD.

7. 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/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>.

Appendix A. Contributors

Shunwan Zhuang Huawei Email: zhuangshunwan@huawei.com

Penghe Tang Huawei Technologies Email: tangpenghe@huawei.com@huawei.com

Authors' Addresses

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

Hang Shi
Huawei
Beiqing Road
Beijing
China
Email: shihang9@huawei.com

Jie Dong (editor)
Huawei
China
Email: jie.dong@huawei.com

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