

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: July 8, 2018

G. Dawra, Ed.
C. Filsfils
Cisco Systems
D. Bernier
Bell Canada
J. Uttaro
AT&T
B. Decraene
Orange
H. Elmalky
Ericsson
X. Xu
Huawei
F. Clad
K. Talaulikar
Cisco Systems
January 4, 2018

BGP Control Plane Extensions for Segment Routing based Service Chaining
[draft-dawra-idr-bgp-sr-service-chaining-02](#)

Abstract

The BGP Control Plane for the SR service-chaining solution is consistent with the BGP Control Plane for the topological Segment Routing Traffic Engineering (SR-TE) solution.

- o BGP Link-State(BGP-LS) address-family/sub-address-family[RFC7752] is used to discover service and topological characteristics from the network.
- o SR-TE policies[I-D.ietf-idr-segment-routing-te-policy] instantiate source-routed policies that may mix service and topological segments.

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 July 8, 2018.

Copyright Notice

Copyright (c) 2018 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 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	2
2.	BGP-LS Extensions for Service Chaining	3
3.	Illustration	6
4.	IANA Considerations	7
4.1.	Service Type Table	7
4.2.	Segment routing function Identifier(SFI)	8
5.	Manageability Considerations	8
6.	Operational Considerations	8
6.1.	Operations	8
7.	Security Considerations	8
8.	Conclusions	8
9.	Acknowledgements	9
10.	References	9
10.1.	Normative References	9
10.2.	Informative References	10
	Authors' Addresses	12

[1.](#) Introduction

Segments are introduced in the SR architecture [[I-D.ietf-spring-segment-routing](#)]. Segment Routing based Service chaining is well described in Section 6 of [[I-D.clad-spring-segment-routing-service-chaining](#)] document with an example network and services.

This document extend the example to add a Segment Routing Controller (SR-C) to the network, for the purpose of service discovery and SR policy instantiation.

Consider the network represented in Figure 1 below where:

- o A and B are two end hosts using IPv4.
- o S1 is an SR-aware firewall Service.
- o S2 is an SR-unaware DPI Service.




Figure 1: Network with Services

SR Controller (SR-C) is connected to Node 1, but may be attached to any node 1-6 in the network.

SR-C is capable of receiving BGP-LS updates to discover topology, and calculating constrained paths between 1 and 6.

However, if SR-C is configured to computation a constrained path from 1 and 6, including a DPI service (i.e., S2) it is not yet possible due to the lack of service distribution. SR-C does not know where a DPI Service is nor the SID for it. It does not know that S2 is a service it needs.

This document proposes an extension to BGP-LS for Service Chaining to distribute the service information to SR-C. There may be other alternate mechanisms to distribute service information to SR-C and are outside of scope of this document. There are no extensions required in SR-TE Policy SAFI.

2. BGP-LS Extensions for Service Chaining

For an attached service, following data needs to be shared with SR-C:

- o Service SID value (e.g. MPLS label or IPv6 address). Service SID MAY only be encoded as LOC:FUNCT, where LOC is the L most significant bits and FUNCT is the 128-L least significant

bits[I-D.filsfils-spring-srv6-network-programming]. ARGs bits, if any, MAY be set to 0 in the advertised service SID.

- o Function Identifier (Static Proxy, Dynamic Proxy, Shared Memory Proxy, Masquerading Proxy, SR Aware Service etc).
- o Service Type (DPI, Firewall, Classifier, LB etc).
- o Traffic Type (IPv4 OR IPv6 OR Ethernet)
- o Opaque Data (Such as brand and version, other extra information)

[[I-D.clad-spring-segment-routing-service-chaining](#)] defines SR-aware and SR-unaware services. This document will reuse these definitions. Per [[RFC7752](#)] Node Attributes are ONLY associated with the Node NLRI. All non-VPN information SHALL be encoded using AFI 16388 / SAFI 71. VPN information SHALL be encoded using AFI 16388 / SAFI 72 with associated RTs.

This document extends SRv6 Node SID TLV [[I-D.dawra-idr-bgppls-srv6-ext](#)] and SR-MPLS SID/Label TLV [[I-D.ietf-idr-bgp-ls-segment-routing-ext](#)] to associate the Service SID Value with Service-related Information using Service Chaining(SC) Sub-TLV.

Function Sub-TLV [[I-D.dawra-idr-bgppls-srv6-ext](#)] of Node SID TLV encodes Identifier(Function ID) along with associated Function Flags.

A Service Chaining (SC) Sub-TLV in Figure 2 is defined as:

+-----+
Type (2 octet)
+-----+
Length (2 octet)
+-----+
Service Type(ST) (2 octet)
+-----+
Flags (1 octet)
+-----+
Traffic Type(1 octet)
+-----+
RESERVED (2 octet)
+-----+

Figure 2: Service Chaining(SC) Sub-TLV

Where:

Type: 16 bit field. TBD

Length: 16 bit field. The total length of the value portion of the TLV.

Service Type(ST): 16bit field. Service Type: categorizes the Service: (such as "Firewall", "Classifier" etc).

Flags: 8 bit field. Bits SHOULD be 0 on transmission and MUST be ignored on reception.

Traffic Type: 8 Bit field. A bit to identify if Service is IPv4 OR IPv6 OR L2 Ethernet Capable. Where:

Bit 0(LSB): Set to 1 if Service is IPv4 Capable

Bit 1: Set to 1 if Service is IPv6 Capable

Bit 2: Set to 1 if Service is Ethernet Capable

RESERVED: 16bit field. SHOULD be 0 on transmission and MUST be ignored on reception.

Service Type(ST) MUST be encoded as part of SC Sub-TLV.

There may be multiple instances of similar Services that needs to be distinguished. For example, firewalls made by different vendors A and B may need to be identified differently because, while they have similar functionality, their behavior is not identical.

In order for SDN Controller to identify the categories of Services and their associated SIDs, this section defines the BGP-LS extensions required to encode these characteristics and other relevant information about these Services.

Another Optional Opaque Metadata(OM) Sub-TLV of Node SID TLV may encode vendor specific information. Multiple of OM Sub-TLVs may be encoded.

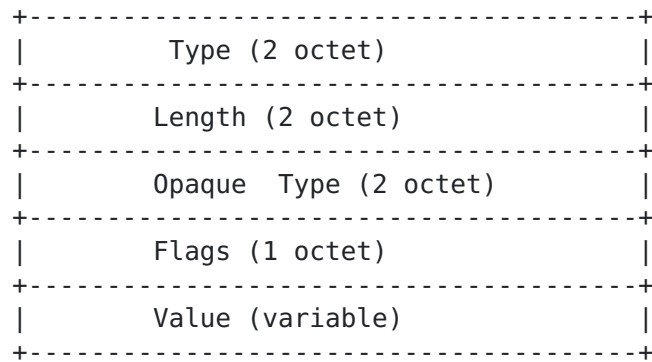


Figure 3: Opaque Metadata(OM) Sub-TLV

- o Type: 16 bit field. TBD.
- o Length: 16 bit field. The total length of the value portion of the TLV.
- o Opaque Type: 8-bit field. Only publishers and consumers of the opaque data are supposed to understand the data.
- o Flags: 8 bit field. Bits SHOULD be 0 on transmission and MUST be ignored on reception.
- o Value: Variable Length. Based on the data being encoded and length is recorded in length field.

Opaque Metadata(OM) Sub-TLV defined in Figure 3 may encode propriety or Service Opaque information such as:

- o Vendor specific Service Information.
- o Traffic Limiting Information to particular Service Type.
- o Opaque Information unique to the Service
- o Propriety Enterprise Service specific Information.

3. Illustration

In our SRv6 example above Figure 1 , Node 5 is configured with an SRv6 dynamic proxy segments (End.AD) C5::AD:F2 for S2.

The BGP-LS advertisement MUST contain and Node SID TLV:

- o Service SID: C5::AD:F2 SID

- o Function ID: END.AD

The BGP-LS advertisement MUST contain a SC Sub-TLV with:

- o Service Type: Deep Packet Inspection(DPI)
- o Traffic Type: IPv4 Capable.

The BGP-LS advertisement MAY contain a OM Sub-TLV with:

- o Opaque Type: Cisco DPI Version
- o Value: 3.5

In our example in Figure 1, using BGP SR-TE SAFI Update [[I-D.ietf-idr-segment-routing-te-policy](#)], SR Controller computes the candidate path and pushes the Policy.

SRv6 encapsulation policy < CF1::, C3::, C5::AD:F2, C6::D4:B > is signaled to Node 1 which has mix of service and topological segments.

[4. IANA Considerations](#)

This document requests assigning code-points from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs".

[4.1. Service Type Table](#)

IANA is request to create a new top-level registry called "Service Type Table (STT)". Valid values are in the range 0 to 65535. Values 0 and 65535 are to be marked "Reserved, not to be allocated".

Service Value(TBD)	Service	Reference	Date
32	Classifier	ref-to-set	date-to-set
33	Firewall	ref-to-set	date-to-set
34	Load Balancer	ref-to-set	date-to-set
35	DPI	ref-to-set	date-to-set

Figure 4

4.2. Segment routing function Identifier(SFI)

IANA is request to extend a top-level registry called "Segment Routing Function Identifier(SFI)" with new code points. This document extends the SFI values defined in [\[I-D.dawra-idr-bgpls-srv6-ext\]](#). Details about the Service functions are defined in [\[I-D.clad-spring-segment-routing-service-chaining\]](#).

Function	Function Identifier
Static Proxy	8
Dynamic Proxy	9
Shared Memory Proxy	10
Masquerading Proxy	11
SRv6 Aware Service	12

5. Manageability Considerations

This section is structured as recommended in [\[RFC5706\]](#)

6. Operational Considerations

6.1. Operations

Existing BGP and BGP-LS operational procedures apply. No additional operation procedures are defined in this document.

7. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the 'Security Considerations' section of [\[RFC4271\]](#) for a discussion of BGP security. Also refer to [\[RFC4272\]](#) and [\[RFC6952\]](#) for analysis of security issues for BGP.

8. Conclusions

This document proposes extensions to the BGP-LS to allow discovery of Services using Segment Routing.

9. Acknowledgements

The authors would like to thank Krishnaswamy Ananthamurthy for his review of this document.

10. References

10.1. Normative References

- [I-D.clad-spring-segment-routing-service-chaining]
Clad, F., Filsfils, C., Camarillo, P.,
daniel.bernier@bell.ca, d., Decraene, B., Peirens, B.,
Yadlapalli, C., Xu, X., Salsano, S., Abdelsalam, A., and
G. Dawra, "Segment Routing for Service Chaining", [draft-clad-spring-segment-routing-service-chaining-00](#) (work in progress), October 2017.
- [I-D.dawra-idr-bgpls-srv6-ext]
Dawra, G., Filsfils, C., Talaulikar, K., Sreekantiah, A.,
and L. Ginsberg, "BGP Link State extensions for IPv6
Segment Routing(SRv6)", [draft-dawra-idr-bgpls-srv6-ext-00](#)
(work in progress), October 2017.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis",
[RFC 4272](#), DOI 10.17487/RFC4272, January 2006,
<<https://www.rfc-editor.org/info/rfc4272>>.
- [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>>.
- [RFC5706] Harrington, D., "Guidelines for Considering Operations and
Management of New Protocols and Protocol Extensions",
[RFC 5706](#), DOI 10.17487/RFC5706, November 2009,
<<https://www.rfc-editor.org/info/rfc5706>>.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of
BGP, LDP, PCEP, and MSDP Issues According to the Keying
and Authentication for Routing Protocols (KARP) Design
Guide", [RFC 6952](#), DOI 10.17487/RFC6952, May 2013,
<<https://www.rfc-editor.org/info/rfc6952>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and
S. Ray, "North-Bound Distribution of Link-State and
Traffic Engineering (TE) Information Using BGP", [RFC 7752](#),
DOI 10.17487/RFC7752, March 2016,
<<https://www.rfc-editor.org/info/rfc7752>>.

10.2. Informative References

[I-D.dawra-bgp-srv6-vpn]

(Unknown), (., Dawra, G., Filsfils, C., Dukes, D., Brissette, P., Camarillo, P., Leddy, J., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Steinberg, D., Raszuk, R., Decraene, B., and S. Matsushima, "BGP Signaling of IPv6-Segment-Routing-based VPN Networks", [draft-dawra-bgp-srv6-vpn-00](#) (work in progress), March 2017.

[I-D.filsfils-spring-segment-routing-policy]

Filsfils, C., Sivabalan, S., Raza, K., Liste, J., Clad, F., Talaulikar, K., Hegde, S., daniel.voyer@bell.ca, d., Lin, S., bogdanov@google.com, b., Horneffer, M., Steinberg, D., Decraene, B., Litkowski, S., and P. Mattes, "Segment Routing Policy for Traffic Engineering", [draft-filsfils-spring-segment-routing-policy-04](#) (work in progress), December 2017.

[I-D.filsfils-spring-srv6-network-programming]

Filsfils, C., Leddy, J., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Steinberg, D., Raszuk, R., Matsushima, S., Lebrun, D., Decraene, B., Peirens, B., Salsano, S., Naik, G., Elmalky, H., Jonnalagadda, P., Sharif, M., Ayyangar, A., Mynam, S., Henderickx, W., Bashandy, A., Raza, K., Dukes, D., Clad, F., and P. Camarillo, "SRv6 Network Programming", [draft-filsfils-spring-srv6-network-programming-03](#) (work in progress), December 2017.

[I-D.ietf-6man-segment-routing-header]

Previdi, S., Filsfils, C., Raza, K., Leddy, J., Field, B., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Matsushima, S., Leung, I., Linkova, J., Aries, E., Kosugi, T., Vyncke, E., Lebrun, D., Steinberg, D., and R. Raszuk, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-07](#) (work in progress), July 2017.

[I-D.ietf-bess-evpn-prefix-advertisement]

Rabadan, J., Henderickx, W., Drake, J., Lin, W., and A. Sajassi, "IP Prefix Advertisement in EVPN", [draft-ietf-bess-evpn-prefix-advertisement-09](#) (work in progress), November 2017.

[I-D.ietf-idr-bgp-ls-segment-routing-ext]

Previdi, S., Psenak, P., Filsfils, C., Gredler, H., and M. Chen, "BGP Link-State extensions for Segment Routing", [draft-ietf-idr-bgp-ls-segment-routing-ext-03](#) (work in progress), July 2017.

[I-D.ietf-idr-bgp-prefix-sid]

Previdi, S., Filsfils, C., Lindem, A., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix SID extensions for BGP", [draft-ietf-idr-bgp-prefix-sid-08](#) (work in progress), January 2018.

[I-D.ietf-idr-segment-routing-te-policy]

Previdi, S., Filsfils, C., Mattes, P., Rosen, E., and S. Lin, "Advertising Segment Routing Policies in BGP", [draft-ietf-idr-segment-routing-te-policy-01](#) (work in progress), December 2017.

[I-D.ietf-isis-segment-routing-extensions]

Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", [draft-ietf-isis-segment-routing-extensions-15](#) (work in progress), December 2017.

[I-D.ietf-spring-segment-routing]

Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-14](#) (work in progress), December 2017.

[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>>.

[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>>.

[RFC5549] Le Faucheur, F. and E. Rosen, "Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop", [RFC 5549](https://www.rfc-editor.org/info/rfc5549), DOI 10.17487/RFC5549, May 2009, <<https://www.rfc-editor.org/info/rfc5549>>.

Authors' Addresses

Gaurav Dawra (editor)
Cisco Systems
USA

Email: gdawra.ietf@gmail.com

Clarence Filsfils
Cisco Systems
Belgium

Email: cfilsfil@cisco.com

Daniel Bernier
Bell Canada
Canada

Email: daniel.bernier@bell.ca

Jim Uttaro
AT&T
USA

Email: ju1738@att.com

Bruno Decraene
Orange
France

Email: bruno.decraene@orange.com

Hani Elmalky
Ericsson
USA

Email: hani.elmalky@gmail.com

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

Francois Clad
Cisco Systems
France

Email: fclad@cisco.com

Ketan Talaulikar
Cisco Systems
India

Email: ketant@cisco.com