

BESS Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 8, 2021

J. Drake
Juniper Networks
A. Farrel
Old Dog Consulting
L. Jalil
Verizon
A. Lingala
AT&T
February 4, 2021

BGP-LS Filters : A Framework for Network Slicing and Enhanced VPNs
draft-drake-bess-enhanced-vpn-06

Abstract

Future networks that support advanced services, such as those enabled by 5G mobile networks, envision a set of overlay networks each with different performance and scaling properties. These overlays are known as network slices and are realized over a common underlay network. In the context of IETF technologies, they are known as IETF network slices.

In order to support IETF network slicing, as well as to offer enhanced VPN services in general, it is necessary to define a mechanism by which specific resources (links and/or nodes) of an underlay network can be used by a specific network slice, VPN, or set of VPNs. This document sets out such a mechanism for use in Segment Routing networks.

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 August 8, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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	3
2.	Requirements Language	4
3.	Overview of Approach	4
4.	Detailed Protocol Operation	6
4.1.	The BGP-LS Filter Attribute	8
4.1.1.	The Filter TLV	9
4.1.2.	The DSCP List TLV	10
4.1.3.	The Color List TLV	11
4.1.4.	The Root TLV	12
4.2.	Error Handling	13
5.	Comparison With ACTN	13
6.	Examples	14
6.1.	MP2MP Connectivity	15
6.2.	P2MP Unidirectional Connectivity	15
6.3.	P2P Unidirectional Connectivity	16
6.4.	P2P Bidirectional Connectivity	17
7.	Security Considerations	18
8.	Manageability Considerations	18
9.	IANA Considerations	19
9.1.	New BGP Path Attribute	19
9.2.	New BGP-LS Filter attribute TLVs Type Registry	19
10.	Acknowledgements	20
11.	Contributors	20
12.	References	20
12.1.	Normative References	20
12.2.	Informative References	21
	Authors' Addresses	23

1. Introduction

Network slicing is an approach to network operations that builds on the concept of network abstraction to provide programmability, flexibility, and modularity. Driven largely by needs surfacing from 5G, the concept of network slicing has gained traction, for example in [\[TS23501\]](#) and [\[TS28530\]](#). Network slicing requires the underlying network to support partitioning the network resources to provide the client with dedicated (private) networking, computing, and storage resources drawn from a shared pool. The network slices may be seen as (and operated as) virtual networks. In the context of IETF technologies network slices are known as "IETF network slices" [\[I-D.ietf-teas-ietf-network-slice-definition\]](#), however, in this document we simply use the term "network slice" since we are working entirely within this context.

Advanced services drive a need to create virtual networks with enhanced characteristics. The tenant of such a virtual network can require a degree of isolation and performance that previously could only be satisfied by dedicated networks. Additionally, the tenant may ask for some level of control to their virtual networks, e.g., to customize the service forwarding paths in the underlying network.

The concept of "IETF network slices" is introduced in [\[I-D.ietf-teas-ietf-network-slice-definition\]](#) and [\[I-D.nsdt-teas-ns-framework\]](#). [\[I-D.ietf-teas-enhanced-vpn\]](#) builds on this concept and introduces "enhanced VPNs".

In order to support network slicing, as well as to offer enhanced VPN services in general, it is necessary to define a mechanism by which specific resources (links and/or nodes) of an underlay network can be used by a specific network slice, a single VPN, or a well-defined set of VPNs. This document sets out such a mechanism for use in Segment Routing networks [\[RFC8402\]](#) and builds on the ideas introduced in [\[I-D.ietf-idr-segment-routing-te-policy\]](#). I.e., it generalizes that work to support multipoint-to-multipoint (MP2MP), point-to-multipoint (P2MP), and bidirectional point-to-point (P2P) topologies; it integrates BGP-based VPN support ([\[RFC4364\]](#), [\[RFC7432\]](#)); it supports Differentiated Services Code Points (DSCP) as well a Color-based forwarding, and it uses BGP Link-State (BGP-LS) [\[RFC7752\]](#) to distribute topology information.

This document supports the concept of a network slice network model interface that provides the function needed by the network slice service model interface defined in [\[I-D.nsdt-teas-ns-framework\]](#).

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.

3. Overview of Approach

The approach described in this document is based on a network controller that uses the {source, destination} traffic matrix and the performance and scaling properties of each network slice, VPN, or set of VPNs in conjunction with the topology of the underlay network to assign each network slice, VPN, or set of VPNs a set of underlay links and nodes that it can use. That is, each network slice, VPN, or set of VPNs gets a subset, either dedicated or shared, of the resources in the underlay network. Note that, in this document, we recognize that scalability of protocol mechanisms to partition network resources is very important; this gives rise to the concept of "a set of VPNs" so that the slice of network resources achieved using the protocol mechanisms defined in this document can be shared by a well-defined set of VPNs as configured by the network operator.

It should be noted that resources can be assigned at any of the following granularities:

- o All provider edge (PE) routers in a given VPN.
- o A set of PEs in a given VPN.
- o An individual PE in a given VPN.

There are two phases to this approach:

Step-1: Discovery and data gathering. Information is gathered from the underlay network about the links, nodes, and network resources available for use by the VPN or network slice.

Step-2: Configuration and provisioning. The underlay resources are configured for use for the VPN or network slice.

Once the network controller has determined the resource assignments, it distributes this information to the PEs that participate in each VPN using the usual VPN information dissemination tools, e.g., route targets (RT) [[RFC4360](#)], route reflectors (RR) [[RFC4456](#)], and RT constraints [[RFC4684](#)].

This information is distributed to the PEs by giving them a customized and limited view of the underlay network on the basis of a network slice, a VPN, or a set of VPNs. Each PE will have a complete view of the underlay network and this customized and limited view acts as filter on the underlay network telling the PE which underlay network resources it can use to direct the traffic of a given network slice, VPN, or set of VPNs to best deliver end-to-end services.

The resource allocation information is encoded using BGP-LS. This approach is chosen for the following reasons:

- o It is BGP-based so it integrates easily with the existing BGP-based VPN infrastructure ([[RFC4364](#)], [[RFC4684](#)]).
- o It supports Segment Routing which is necessary to enforce the PEs' usage of the resources allocated to the VPN or set of VPNs.
- o It supports Segment Routing which is necessary to enforce the PEs' usage of the resources allocated to the network slice, VPN, or set of VPNs. The use of RSVP-TE ([[RFC3209](#)]) rather than Segment Routing is at the discretion of the network operator as BGP-LS supports both and either confines a packet flow to a specific path.
- o It supports inter-AS connectivity which is a prerequisite for supporting the existing BGP-based VPN infrastructure.
- o It is canonical, in that it can be used to advertise the resources of underlay networks that use either IS-IS or OSPF.

It should be noted that this mechanism also follows the scalability model of the existing BGP-based VPN infrastructure, which is that the per-VPN information is restricted to only those PE routers that are supporting that VPN and that the provider (P) routers have no per-VPN state.

The PEs in non-enhanced VPNs do not receive this resource allocation information and would not confine their usage of the underlay network resources. In order to ensure that the underlay network resources allocated to enhanced VPNs are not inadvertently used by the PEs in non-enhanced VPNs, the network controller SHOULD ensure that the IGP and traffic engineering (TE) metrics for these resources is higher than the metrics for the underlay network resources allocated to non-enhanced VPNs. In certain situations, detailed in [Section 4](#), PEs in enhanced VPNs will use the underlay networks resources allocated to non-enhanced VPNs.

Additional to the programming of the PEs and its computation and assignment of resources for use by network slices, VPNs, or sets of VPNs, the network controller also instructs the P routers to make the actual allocation of these resources by assigning link bandwidth to a specific DSCP or adjacency segment identifier (SID) [[I-D.dong-spring-sr-for-enhanced-vpn](#)].

4. Detailed Protocol Operation

We define a BGP-LS Filter to be a BGP-LS encoded description of a subset of the links and nodes in the underlay network. A BGP-LS Filter defines all or part of the topology for a network slice or a set of one or more VPNs. The topology defined by a BGP-LS Filter needs to provide connectivity between the PEs in a given network slice, VPN or set of VPNs. I.e., it connects the PEs in these VPNs and is used by them to send packets to each other. A given filter is tagged with the route targets of the VPNs whose PEs are to import the filter. A BGP-LS Filter is pushed southbound to those PEs by the network controller and SHOULD provide multiple paths between a given ingress/egress PE pair.

Note that there will be multiple BGP-LS Filters in a given network deployment and that a given underlay network link or node may appear in more than one of them. In order to provide disambiguation, the address family indicator (AFI) 16388 (BGP-LS) and the subsequent address family identifier (SAFI) 72 (BGP-LS-VPN) are used in BGP-LS UPDATE messages and the network controller SHOULD allocate a different route distinguisher (RD) to each BGP-LS Filter. As for standard VPNs, an implementation option ("RD Auto") may be offered to assist in configuring unique RDs.

Within a given VPN, when an ingress PE needs to send a packet to an egress PE it selects a path to that egress PE from the topology defined by the BGP-LS Filters it has imported for that VPN. It then either adds a segment routing label stack specifying that path to the packet or places the packet in an RSVP-TE LSP which uses that path. The ingress PE may use any path computation it wishes if that path computation confines the path to the topology defined by the relevant set of BGP-LS Filters.

If Segment Routing is used and a node SID or a prefix SID is placed in the segment routing label stack, then when that segment is active the P routers will forward the packet using the underlay network resources allocated to non-enhanced VPNs. Similarly, if the RSVP-TE label switched path (LSP) was established using a loose source route to the subject node, the path to that node was selected using the underlay network resources allocated to non-enhanced VPNs.

Because the BGP-LS UPDATE messages specifying a BGP-LS Filter may arrive in any order and the BGP-LS UPDATE messages of multiple BGP-LS Filters may be interleaved, there is a need for a new attribute that is attached to a BGP-LS UPDATE. This attribute contains a Filter ID, a Filter version number, a Filter type (MP2MP, P2MP, or P2P), the total number of fragments in the filter, and the specific fragment number of the piece in hand. I.e., it is assumed that a PE may import more than one BGP-LS Filter, that a given BGP-LS Filter may change over time, and that a given BGP-LS Filter may span multiple BGP-LS UPDATE messages. The Filter ID needs to be unique across the set of VPNs into which the BGP-LS Filter is to be imported.

A BGP-LS Filter that is created for a set of VPNs will contain a set of network resources sufficient to connect between the PEs in each discrete VPN in the set, and each of the BGP-LS UPDATE messages for the filter MUST be tagged with the RT for each VPN in the set.

If a PE imports more than one BGP-LS Filter it MAY use the union of the links and nodes specified in each filter when selecting a path.

A given BGP-LS Filter may change in response to updates to the PE membership in a VPN to which the BGP-LS Filter applies or to updates to the underlay network. This implies that the network controller needs to be connected to the route reflectors associated with the VPNs for which it is providing BGP-LS maps. When this occurs, the network controller SHOULD push a new version of the affected BGP-LS Filters. That is, it increments the version number of each BGP-LS Filter. Note that a network controller does not need to compute new BGP-LS Filters in response to an individual link or node failure in the underlay network if connectivity still exists among the PEs in the network slice, VPN or set of VPNs with the existing BGP-LS Filters.

A BGP-LS Filter cannot be used by a PE until it is completely assembled. If the BGP-LS Filter that is being assembled is a newer version of a BGP-LS Filter that the PE is currently using, the PE SHOULD continue to use its current version of the BGP-LS Filter until the newer version is completely assembled.

When selecting a path using one or more BGP-LS Filters, an ingress PE can use a link or node only if it is active in the underlay network. If this precludes connectivity to the egress PE it may use the underlay network resources not allocated to enhanced VPNs to reach the egress PE.

Additionally, when there is a newly activated PE it will not be present in any of the BGP-LS Filters used by the other PEs. Until a new BGP-LS Filter that contains that PE has been distributed, other

PEs will use the underlay network resources not allocated to enhanced VPNs to reach the newly activated PE, and the newly activate PE will use these resources to reach other PEs.

4.1. The BGP-LS Filter Attribute

[RFC4271] defines the BGP Path attribute. This document introduces a new Optional Transitive Path attribute called the BGP-LS Filter attribute with value TBD1 to be assigned by IANA.

The first BGP-LS Filter attribute MUST be processed and subsequent instances MUST be ignored.

The common fields of the BGP-LS Filter attribute are set as follows:

- o Optional bit is set to 1 to indicate that this is an optional attribute.
- o The Transitive bit is set to 1 to indicate that this is a transitive attribute.
- o The Extended Length bit is set according to the length of the BGP-LS Filter attribute as defined in [[RFC4271](#)].
- o The Attribute Type Code is set to TBD1.

The content of the BGP-LS Filter attribute is a series of Type-Length-Value (TLV) constructs. Each TLV may include sub-TLVs. All TLVs and sub-TLVs have a common format that is:

- o Type: A single octet indicating the type of the BGP-LS Filter attribute TLV. Values are taken from the registry described in [Section 9.2](#).
- o Length: A two octet field indicating the length of the data following the Length field counted in octets.
- o Value: The contents of the TLV.

The formats of the TLVs defined in this document are shown in the following sections. The presence rules and meanings are as follows.

- o The BGP-LS Filter attribute MUST contain a Filter TLV.
- o The BGP-LS Filter attribute MAY contain a DSCP List TLV.
- o The BGP-LS Filter attribute MAY contain a Color List TLV.

- o The BGP-LS Filter attribute MAY contain a Root TLV.

4.1.1. The Filter TLV

The BGP-LS Filter attribute MUST contain exactly one Filter TLV. Its format is shown in Figure 1. Note that a given BGP-LS Filter may span multiple UPDATE messages and the Topology, Version Number, and the Number of Fragments fields in the BGP-LS Filter attribute contained in each UPDATE message MUST be set to the same value or the BGP-LS Filter is unusable.

```

+-----+
|   Type = 1 (1 octet)   |
+-----+
|   Length (2 octets)   |
+-----+
|   Topology (1 Octet)  |
+-----+
|   ID (4 Octets)       |
+-----+
|   Version Number (4 Octets)
+-----+
|   Number of Fragments (4 Octets)
+-----+
|   Fragment Number (4 Octets)
+-----+

```

Figure 1: The Filter TLV Format

The fields are as follows:

- o Type is set to 1 to indicate a Filter TLV.
- o Length is set to 17 octets.
- o Topology indicates the topology defined by this BGP-LS Filter.
 1. P2P unidirectional
 2. P2P bidirectional
 3. P2MP
 4. MP2MP

- o The ID of this BGP-LS Filter. This ID needs to be unique within the set of VPNs into which the BGP-LS Filter is to be imported.
- o The Version Number of this BGP-LS Filter. The contents of a BGP-LS Filter with a given ID may change over time. This field indicates the version of the BGP-LS Filter being advertized in this UPDATE message.
- o Number of Fragments indicates the number of BGP UPDATE messages defining this BGP-LS Filter.
- o Fragment Number indicates ordinal position of this UPDATE message within the set of UPDATE messages defining this BGP-LS Filter. A BGP-LS Filter is not complete, i.e., usable, until all UPDATE messages have been received with Fragment Numbers in the range $1 \leq \text{Fragment Number} \leq \text{Number of Fragments}$. An UPDATE message with a Fragment Number outside this range is to be ignored.

4.1.2. The DSCP List TLV

The DSCP List TLV MAY be included in the BGP-LS Filter attribute. If included, a packet whose DSCP matches a DSCP in the DSCP list is to be forwarded using the BGP-LS Filter defined by the BGP-LS Filter attribute that contains the DSCP list. The first DSCP List TLV MUST be processed and subsequent instances MUST be ignored. The format of the DSCP List TLV is shown in Figure 2.

If a DSCP List TLV is included in a BGP-LS Filter attribute, then a packet that matches an entry in the list MAY be forwarded using the BGP-LS Filter defined by the BGP-LS Filter attribute, but a packet which doesn't match an entry in this list MUST NOT use the filter. If both a DSCP List TLV and a Color List TLV (see [Section 4.1.3](#)) are both included in a BGP-LS Filter attribute, packets matching an entry in either list MAY be forwarded using the BGP-LS Filter defined by the BGP-LS Filter attribute. If neither list is included in a BGP-LS Filter attribute, then all packets for that network slice, VPN, or set of VPNs can be forwarded using the BGP-LS Filter defined by the containing BGP-LS Filter attribute.




Figure 2: The DSCP List TLV Format

The fields are as follows:

- o Type is set to 2 to indicate a DSCP List TLV.
- o Length indicates the length in octets of the DSCP List.
- o DSCP List contains a list of DSCPs, each one octet in length and encodes the DSCP per [\[RFC2474\]](#) as the most significant six bits of the octet.

4.1.3. The Color List TLV

The Color List TLV MAY be included in the BGP-LS Filter attribute. If a BGP UPDATE contains a Color extended community with a color (as defined by [\[I-D.ietf-idr-tunnel-encaps\]](#)) that matches an entry in the Color List, then a packet whose destination is covered by one of the routes in that UPDATE is to be forwarded using the BGP-LS Filter defined by the BGP-LS Filter attribute that contains the Color List TLV. The first Color List TLV MUST be processed and subsequent instances MUST be ignored. The format of the Color List TLV is shown in Figure 3.

If Color List TLV is included in a BGP-LS Filter attribute, then a packet that matches an entry in the list MAY be forwarded using the BGP-LS Filter defined by the BGP-LS Filter attribute, but a packet which doesn't match an entry in this list MUST NOT use the filter. If both a DSCP List TLV (see [Section 4.1.2](#) and a Color List TLV are both included in a BGP-LS Filter attribute, packets matching an entry in either list MAY be forwarded using the BGP-LS Filter defined by the BGP-LS Filter attribute. If neither list is included in a BGP-LS Filter attribute, then all packets for that network slice, VPN, or set of VPNs can be forwarded using the BGP-LS Filter defined by the containing BGP-LS Filter attribute.




Figure 3: The Color List TLV Format

The fields are as follows:

- o Type is set to 3 to indicate a Color List TLV.
- o Length indicates the length in octets of the Color List.
- o Color List contains a list of Colors, each four octets in length and as defined in [[I-D.ietf-idr-tunnel-encaps](#)].

4.1.4. The Root TLV

The Root TLV MUST be included in the BGP-LS Filter attribute if its topology is of type P2MP or P2P unidirectional. It defines the root node for that topology and if it is not present the BGP-LS Filter is unusable. The TLV, if present, MUST be ignored if the topology is of type MP2MP or P2P bidirectional.

The Root TLV is structured as shown in Figure 4 and MAY contain any of the sub-TLVs defined in [section 3.2.1.4 of \[RFC7752\]](#).




Figure 4: The Root TLV Format

The fields are as follows:

- o Type is set to 3 to indicate a Color List TLV.
- o Length indicates the length in octets of the Color List.

- o There follows a sequence of zero or more sub-TLVs as defined in [section 3.2.1.4 of \[RFC7752\]](#). The presence of sub-TLVs can be deduced from the Length field of the Root TLV.

4.2. Error Handling

[Section 6 of \[RFC4271\]](#) describes the handling of malformed BGP attributes, or those that are in error in some way. [\[RFC7606\]](#) revises BGP error handling specifically for the for UPDATE message, provides guidelines for the authors of documents defining new attributes, and revises the error handling procedures for a number of existing attributes. This document introduces the BGP-LS Filter attribute and so defines error handling as follows:

- o When parsing a message, an unknown Attribute Type code or a length that suggests that the attribute is longer than the remaining message is treated as a malformed message and the "treat-as-withdraw" approach is used as per [\[RFC7606\]](#).
- o When parsing a message that contains an BGP-LS Filter attribute, the following cases constitute errors:
 1. Optional bit is set to 0 in BGP-LS Filter attribute.
 2. Transitive bit is set to 0 in BGP-LS Filter attribute.
 3. The attribute does not contain a Filter TLV or contains more than one Filter TLV.
 4. The TLV length indicates that the TLV extends beyond the end of the BGP-LS Filter attribute.
 5. There is an unknown TLV type field found in BGP-LS Filter attribute.

The errors listed above are treated as follows:

- 1., 2., 3., 4.: The attribute MUST be treated as malformed and the "treat-as-withdraw" approach used as per [\[RFC7606\]](#).
- 5.: Unknown TLVs SHOULD be ignored, and message processing SHOULD continue.

5. Comparison With ACTN

Abstraction and Control of TE Networks (ACTN) [\[RFC8453\]](#) is a framework that facilitates the abstraction of underlying network resources to higher-layer applications and that allows network

operators to create virtual networks through the abstraction of the operators' network resources. The applicability of ACTN to network slicing is discussed further in [\[I-D.king-teas-applicability-actn-slicing\]](#).

Essentially the ACTN framework describes how to request and provision a network slice, but does not define how the network is operated to deliver that slice. Therefore, a direct comparison between this work and ACTN is not appropriate. ACTN could be used as a management framework to operate a slicing system built using the protocol extensions defined in this document.

6. Examples

Figure 5 shows a sample underlay topology. Six PEs (PE1 through PE6) are connected across a network of twelve P nodes (P1 through P12). Each PE is dual-homed, and the P nodes are variously connected so that there are multiple routes between PEs.




Figure 5: Underlay Network Topology

6.1. MP2MP Connectivity

Figure 6 shows how a Multi-point-to-multipoint (MP2MP) service that connects PE1, PE3, and PE6 can be installed over the underlay network. Paths have been computed so that, for example, PE1 is connected to both PE3 and PE6 via pairs of redundant paths. Similarly, PE3 is connected to PE1 and PE6, and PE6 is connected to PE1 and PE3.




Figure 6: An MP2MP Service Installed at PE1, PE3, and PE6

6.2. P2MP Unidirectional Connectivity

Figure 7 shows the provision of a Point-to-Multipoint (P2MP) service rooted at PE3 and connected to PE1 and PE6. As in the previous example, a pair of redundant paths is established between PE3 and each of PE1 and PE6. Thus, the two paths from PE3 to PE1 are PE3-P1-P4-P7-PE1 and PE3-P2-P9-P8-PE1.




Figure 7: A P2MP Unidirectional Service Installed at PE3

6.3. P2P Unidirectional Connectivity

Figure 8 shows a Point-to-Point (P2P) service rooted at PE1 and connected to PE3. This is equivalent to a Segment Routing Traffic Engineering (SR TE) Policy [[I-D.ietf-idr-segment-routing-te-policy](#)] installed at PE1.

As in the previous examples, a pair of redundant paths are computed.




Figure 8: A P2P Unidirectional Service (SR TE Policy) Installed at PE1

6.4. P2P Bidirectional Connectivity

Figure 9 show a bidirectional P2P service connecting PE1 and PE6. This is equivalent to a Segment Routing Traffic Engineering (SR TE) Policy [[I-D.ietf-idr-segment-routing-te-policy](#)] installed at PE1 and PE6.




Figure 9: A P2P Bidirectional Service Installed at PE1 and PE6

7. Security Considerations

TBD

8. Manageability Considerations

Per VPN OAM and telemetry will be required in order to monitor and verify the performance of network slices. This is particularly important when the performance of a network slice has been committed to a customer through a Service Level Agreement.

As noted in [Section 5](#), ACTN may provide a suitable management model. However, an Enhanced VPN service model may be needed following the concepts described in [\[RFC8309\]](#) and similar in structure to the Layer 3 VPN service model defined in [\[RFC8299\]](#).

Local policy may be used to balance load between BGP-LS filters that are matched by the same flow. It MUST be possible for an operator to query those policies and understand how traffic is being matched to filters. An implementation MAY also make those policies configurable by an operator so that the operator can exert control over how load is balanced (for example, by applying weights to various filters).

9. IANA Considerations

9.1. New BGP Path Attribute

IANA maintains a registry of "Border Gateway Protocol (BGP) Parameters" with a subregistry of "BGP Path Attributes". IANA is requested to assign a new Path attribute called "BGP-LS Filter attribute" (TBD1 in this document) with this document as a reference.

9.2. New BGP-LS Filter attribute TLVs Type Registry

IANA maintains a registry of "Border Gateway Protocol (BGP) Parameters". IANA is request to create a new subregistry called the "BGP-LS Filter attribute TLVs" registry.

Valid values are in the range 0 to 255.

- o Values 0 and 255 are to be marked "Reserved, not to be allocated".
- o Values 1 through 254 are to be assigned according to the "First Come First Served" policy [[RFC8126](#)]

This document should be given as a reference for this registry. The new registry should track:

- o Type
- o Name
- o Reference Document or Contact
- o Registration Date

The registry should initially be populated as follows:

Type	Name	Reference	Date
1	Filter TLV	[This.I-D]	Date-to-be-set
2	DSCP List TLV	[This.I-D]	Date-to-be-set
3	Color List TLV	[This.I-D]	Date-to-be-set
4	Root TLV	[This.I-D]	Date-to-be-set

10. Acknowledgements

The authors are grateful to all those who contributed to the discussions that led to this work: Ron Bonica, Stewart Bryant, Jie Dong, Keyur Patel, Julian Lucek, and Colby Barth.

Stephane Litkowski provided useful review comments.

11. Contributors

The following people contributed text to this document:

Gyan Mishra
Email: hayabusagsm@gmail.com

12. References

12.1. Normative References

- [I-D.ietf-idr-tunnel-encaps]
Patel, K., Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", [draft-ietf-idr-tunnel-encaps-21](#) (work in progress), January 2021.
- [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>>.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", [RFC 2474](#), DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.
- [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>>.
- [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>>.

- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", [RFC 7432](#), DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", [RFC 7606](#), DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 8126](#), DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.

12.2. Informative References

- [I-D.dong-spring-sr-for-enhanced-vpn]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Segment Routing based Virtual Transport Network (VTN) for Enhanced VPN", [draft-dong-spring-sr-for-enhanced-vpn-13](#) (work in progress), January 2021.
- [I-D.ietf-idr-segment-routing-te-policy]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Rosen, E., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", [draft-ietf-idr-segment-routing-te-policy-11](#) (work in progress), November 2020.
- [I-D.ietf-teas-enhanced-vpn]
Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Networks (VPN+) Service", [draft-ietf-teas-enhanced-vpn-06](#) (work in progress), July 2020.

[I-D.ietf-teas-ietf-network-slice-definition]

Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "Definition of IETF Network Slices", [draft-ietf-teas-ietf-network-slice-definition-00](#) (work in progress), January 2021.

[I-D.king-teas-applicability-actn-slicing]

King, D., Drake, J., and H. Zheng, "Applicability of Abstraction and Control of Traffic Engineered Networks (ACTN) to Network Slicing", [draft-king-teas-applicability-actn-slicing-08](#) (work in progress), October 2020.

[I-D.nsdtd-teas-ns-framework]

Gray, E. and J. Drake, "Framework for Transport Network Slices", [draft-nsdt-teas-ns-framework-04](#) (work in progress), July 2020.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.

[RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", [RFC 4456](#), DOI 10.17487/RFC4456, April 2006, <<https://www.rfc-editor.org/info/rfc4456>>.

[RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", [RFC 4684](#), DOI 10.17487/RFC4684, November 2006, <<https://www.rfc-editor.org/info/rfc4684>>.

[RFC8299] Wu, Q., Ed., Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", [RFC 8299](#), DOI 10.17487/RFC8299, January 2018, <<https://www.rfc-editor.org/info/rfc8299>>.

[RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", [RFC 8309](#), DOI 10.17487/RFC8309, January 2018, <<https://www.rfc-editor.org/info/rfc8309>>.

- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", [RFC 8453](#), DOI 10.17487/RFC8453, August 2018, <<https://www.rfc-editor.org/info/rfc8453>>.
- [TS23501] 3GPP, "System architecture for the 5G System (5GS) - 3GPP TS23.501", 2016, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>>.
- [TS28530] 3GPP, "Management and orchestration; Concepts, use cases and requirements - 3GPP TS28.530", 2016, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>>.

Authors' Addresses

John Drake
Juniper Networks

Email: jdrake@juniper.net

Adrian Farrel
Old Dog Consulting

Email: adrian@olddog.co.uk

Luay Jalil
Verizon

Email: luay.jalil@verizon.com

Avinash Lingala
AT&T

Email: ar977m@att.com

