Ospf Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 31, 2015

Q. Liang
J. You
Huawei

September 27, 2014

OSPF Extensions for Flow Specification
draft-liang-ospf-flowspec-extensions-01

Abstract

This document discusses the use cases why OSPF (Open Shortest Path First) distributing flow specification (FlowSpec) routes is necessary. This document also defines a new OSPF FlowSpec Opaque Link State Advertisement (LSA) encoding format that can be used to distribute FlowSpec routes.

For the network only deploying IGP (Interior Gateway Protocol) (e.g. OSPF), it is expected to extend IGP to distribute FlowSpec routes. One advantage is to mitigate the impacts of Denial-of-Service (DoS) attacks.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

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 http://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 March 31, 2015.

Table of Contents

1. Introduction

[RFC5575] defines a new Border Gateway Protocol Network Layer Reachability Information (BGP NLRI) encoding format that can be used to distribute traffic flow specifications. One application of that encoding format is to automate inter-domain coordination of traffic filtering, such as what is required in order to mitigate (distributed) denial-of-service attacks. [RFC5575] allows flow

specifications received from an external autonomous system to be forwarded to a given BGP peer. However, in order to block the attack traffic more effectively, it is better to distribute the BGP FlowSpec routes to the customer network, which is much closer to the attacker.

For the network only deploying IGP (e.g. OSPF), it is expected to extend IGP to distribute FlowSpec routes. This document discusses the use cases why OSPF distributing FlowSpec routes is necessary. This document also defines a new OSPF FlowSpec Opaque Link State Advertisement (LSA) [RFC5250] encoding format that can be used to distribute FlowSpec routes to the edge routers in the customer network. This mechanism can be used to mitigate the impacts of DoS attacks.

2. Terminology

This section contains definitions for terms used frequently throughout this document. However, many additional definitions can be found in [RFC5250] and [RFC5575].

Flow Specification (FlowSpec): A flow specification is an n-tuple consisting of several matching criteria that can be applied to IP traffic, including filters and actions. Each FlowSpec consists of a set of filters and a set of actions.

3. Use Cases for OSPF based FlowSpec Distribution

For the network only deploying IGP (e.g. OSPF), it is expected to extend IGP (OSPF in this document) to distribute FlowSpec routes, because when the FlowSpec routes are installed in the customer network, it would be closer to the attacker than when they are installed in the provider network. Consequently, the attack traffic could be blocked or the suspicious traffic could be limited to a low rate as early as possible.

The following sub-sections discuss the use cases for OSPF based FlowSpec routes distribution.

3.1. BGP/MPLS VPN

[RFC5575] defines a BGP NLRI encoding format to distribute traffic

flow specifications in BGP deployed network. However in the BGP/MPLS VPN scenario, the IGP (e.g. IS-IS, OSPF) is used between PE (Provider Edge) and CE (Customer Edge) for many deployments. In order to distribute the FlowSpec routes to the customer network, the IGP needs to support the FlowSpec route distribution. The FlowSpec routes are usually generated by the traffic analyzer or the traffic policy center in the network. Depending on the location of the traffic analyzer deployment, two different distribution scenarios will be discussed below.

3.1.1. Traffic Analyzer Deployed in Provider Network

The traffic analyzer (also acting as the traffic policy center) could be deployed in the provider network as shown in Figure 1. If the traffic analyzer detects attack traffic from the customer network VPN1, it would generate the FlowSpec routes for preventing DoS attacks. The FlowSpec routes with a route distinguisher information corresponding to VPN1 are distributed from the traffic analyzer to the PE1 which the traffic analyzer is the attached to. If the traffic analyzer is also a BGP speaker, it can distribute the FlowSpec routes based on the BGP [RFC5575]. Then the PE1 distributes the FlowSpec routes further to the PE2. Finally, the FlowSpec routes need to be distributed from the PE2 to the CE2 based on OSPF, i.e. to the customer network VPN1. As the attacker is more likely in the customer network, if the FlowSpec routes installed on the CE2, it could mitigate the impacts of DoS attacks better.

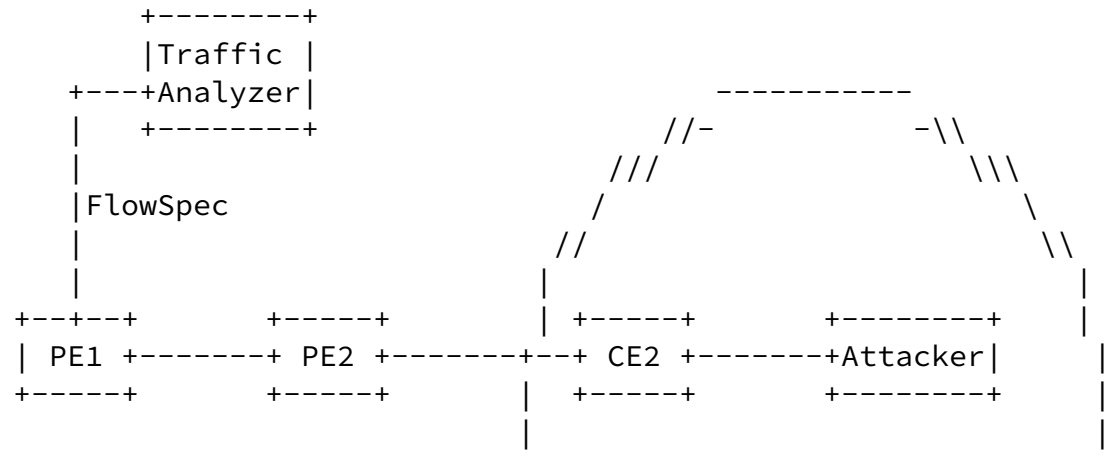

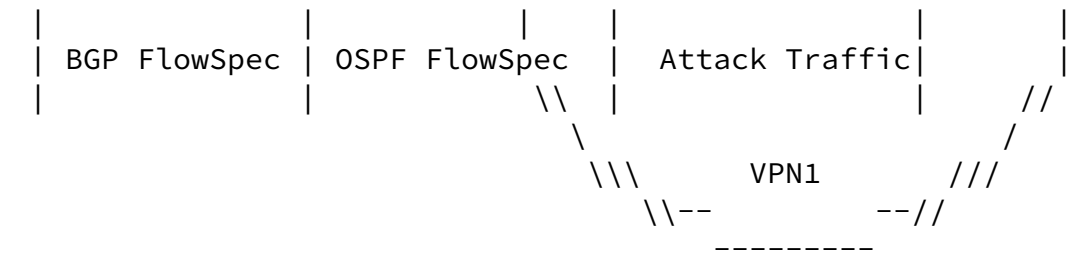

Figure 1: Traffic Analyzer deployed in Provider Network

3.1.2. Traffic Analyzer Deployed in Customer Network

The traffic analyzer (also acting as the traffic policy center) could be deployed in the customer network as shown in Figure 2. If the traffic analyzer detects attack traffic, it would generate FlowSpec routes for preventing DoS attacks. Then the FlowSpec routes would be distributed from the traffic analyzer to the CE1 based on OSPF or other policy protocol (e.g. RESTful API over HTTP). Further, the FlowSpec routes need to be distributed through the provider network

via the PE1/PE2 to the CE2, i.e. to the remote customer network VPN1 Site1. If the FlowSpec routes installed on the CE2, it could block the attack traffic as close to the source of the attack as possible.

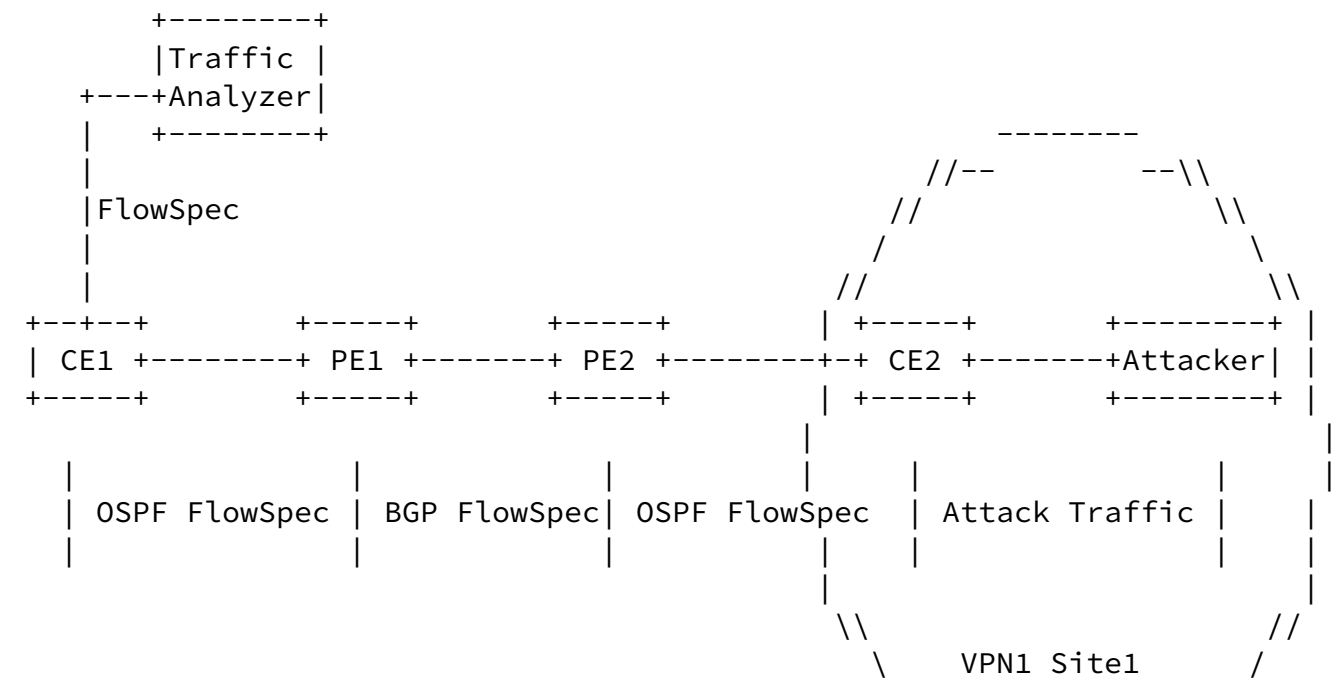

Figure 2: Traffic Analyzer deployed in Customer Network

3.2. OSPF Campus Network

For the network not deploying BGP, for example, the campus network using OSPF, it is expected to extend OSPF to distribute FlowSpec routes as shown in Figure 3. In this kind of network, the traffic analyzer could be deploy with a router, then the FlowSpec routes from the traffic analyzer need to be distributed to the other routers in this domain based on OSPF.

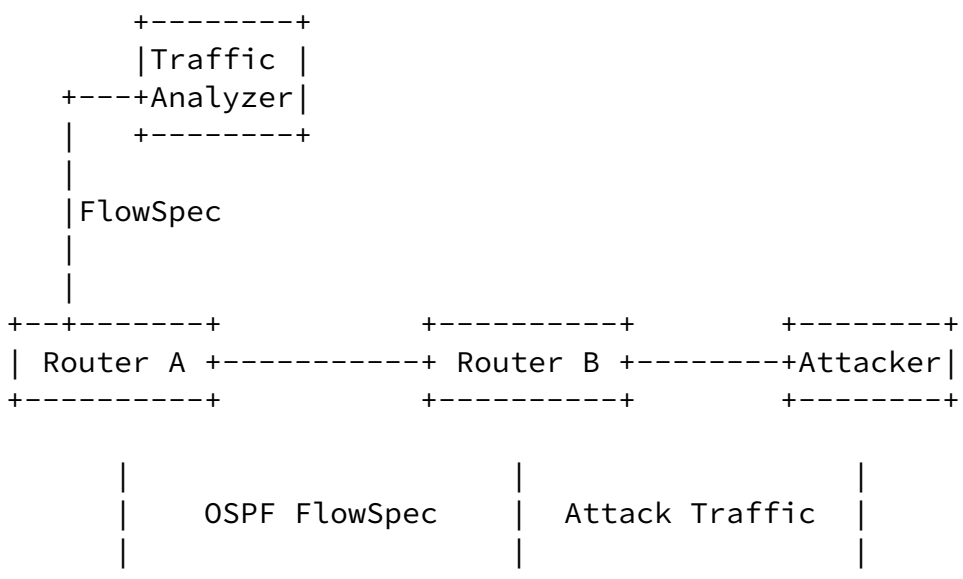

Figure 3: OSPF Campus Network

4. OSPF Extensions for FlowSpec Routes

This document defines a new OSPF flow specification Opaque Link State Advertisement (LSA) encoding format that can be used to distribute traffic flow specifications. This new OSPF FlowSpec Opaque LSA is extended based on [RFC5250].

The FlowSpec Opaque LSA is defined below in Figure 4:

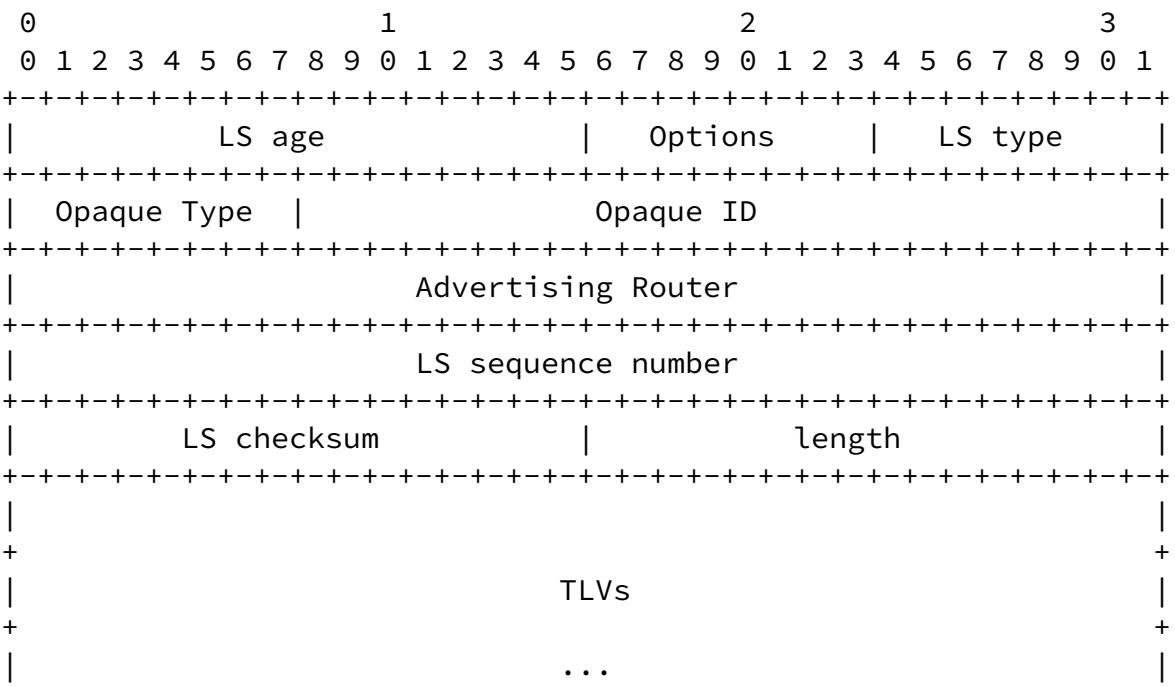

Figure 4: FlowSpec Opaque LSA

LS age: the same as defined in [RFC2328].

Options: the same as defined in [RFC2328].

Link-State Type: A value of 11 denotes that the LSA is flooded throughout the Autonomous System (e.g., has the same scope as type-5 LSAs). Opaque LSAs with AS-wide scope MUST NOT be flooded into stub areas or NSSAs (Not-So-Stubby Areas) [RFC5250].

Opaque type: OSPF FlowSpec Opaque LSA (Type Code: TBD1).

Opaque ID: the same as defined in [RFC5250].

Advertising Router: the same as defined in [RFC2328].

LS sequence number: the same as defined in [RFC2328].

LS checksum: the same as defined in [RFC2328].

Length: the same as defined in [RFC2328].

TLVs: one or more TLVs MAY be included in a FlowSpec Opaque LSA to carry FlowSpec information.

The variable TLVs section consists of one or more nested Type/Length/Value (TLV) tuples. Nested TLVs are also referred to as sub-TLVs. The format of each TLV is shown in Figure 5:

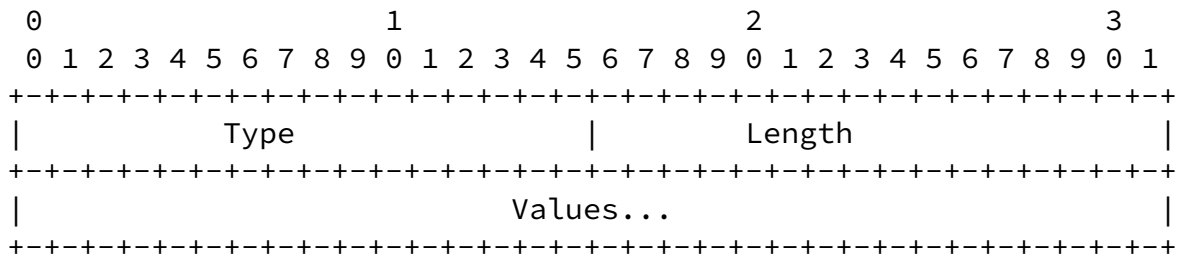

Figure 5: TLV Format

The Length field defines the length of the value portion in octets (thus a TLV with no value portion would have a length of 0). The TLV is padded to 4-octet alignment; padding is not included in the length field (so a 3-octet value would have a length of 3, but the total size of the TLV would be 8 octets). Nested TLVs are also 32-bit aligned. For example, a 1-byte value would have the length field set to 1, and 3 octets of padding would be added to the end of the value portion of the TLV.

4.1. OSPF FlowSpec Filters TLV

The FlowSpec Opaque LSA carries one or more FlowSpec Filters TLVs and corresponding FlowSpec Action TLVs. OSPF FlowSpec Filters TLV is defined below in Figure 6.

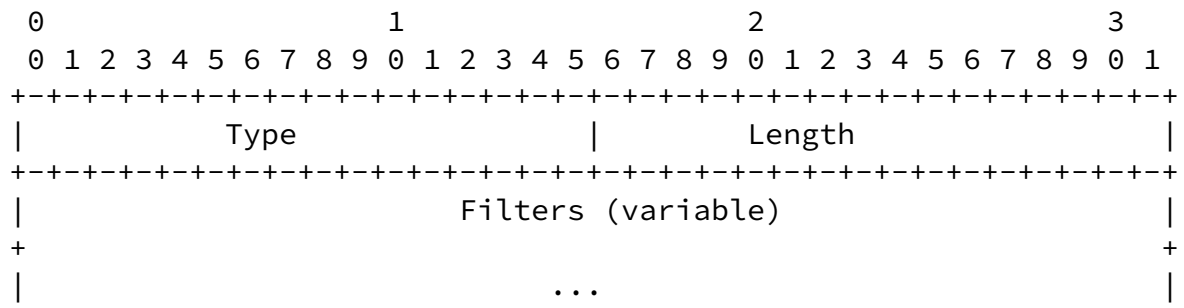

Figure 6: OSPF FlowSpec Filters TLV

Type: the TLV type (Type Code: TBD2)

Length: the size of the value field (typically in bytes)

Filters: the same as "flow-spec NLRI value" defined in [RFC5575].

4.2. OSPF FlowSpec Action TLV

There are one or more FlowSpec Action TLVs associated with a FlowSpec Filters TLV. Meanwhile, different FlowSpec Filters TLV could have the same FlowSpec Action TLV/s. The following OSPF FlowSpec action TLVs are the same as defined in [RFC5575].

Table 1: Traffic Filtering Actions in [RFC5575]

type	FlowSpec Action	encoding
0x8006	traffic-rate	2-byte as#, 4-byte float
0x8007	traffic-action	bitmask
0x8008	redirect	6-byte Route Target
0x8009	traffic-marking	DSCP value

4.3. Capability Advertisement

OSPF routers may use Router Information (RI) LSA [RFC4970] for OSPF features advertisement and discovery. The FlowSpec route requires an additional capability for the OSPF router. This capability needs to be advertised to other routers in an AS. This FlowSpec capability could be advertised in a RI Opaque LSA [RFC4970].

The format of the OSPF Router Informational Capabilities TLV within the body of an RI LSA is defined as follows:

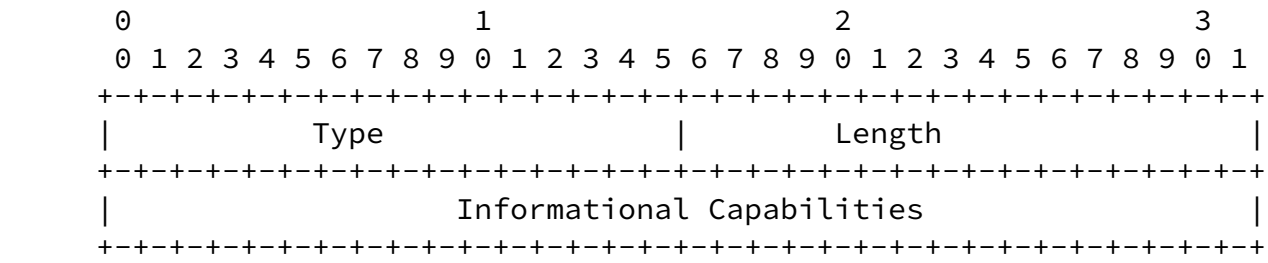

Figure 4: OSPF RI Capabilities TLV

The following informational capability bits are assigned:

Bit	Capabilities
6 (TBD3)	OSPF FlowSpec
7-31	Unassigned (Standards Action)

4.4. Import-policy Extended Community

When FlowSpec routes are from the BGP protocol, these FlowSpec routes need to be imported to the IGP protocol. This document defines a new filtering policy that it standardizes as a BGP extended community value [RFC4360]. This extended community is used to specify a particular action, i.e. importing the FlowSpec routes to the IGP protocol. Thus a new extended community attribute, i.e. import-policy (Type Code: TBD4) is defined as follows:

type	extended community	encoding
TBD4	import-policy	IGP target

The format of the import-policy extended community is defined as follows.

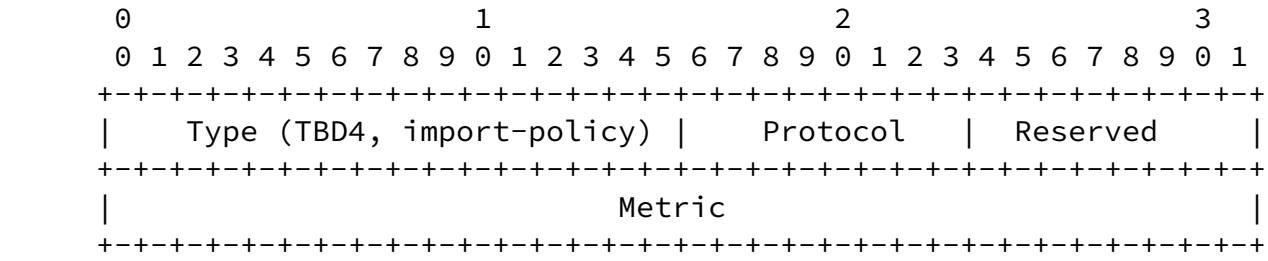

Figure 5: Import-policy Extended Community

This import-policy extended community is with Type Field composed of 2 octets and Value Field composed of 6 octets. The Value Field consists of two sub-fields:

Protocol: 1 octet, this sub-field defines the IGP Type, e.g. 1 for OSPF, 2 for ISIS.

Metric: 4 octets, this sub-field represents the aggregate IGP or TE path cost.

If this import-policy extended community is not present, BGP FlowSpec routes should not be imported to the IGP FlowSpec routing table.

5. IANA Considerations

This document defines a new OSPF Opaque LSA, i.e. FlowSpec Opaque LSA (Type Code: TBD1), which is used to distribute traffic flow specifications.

This document defines OSPF FlowSpec Filters TLV (Type Code: TBD2), which is used to describe the filters.

This document defines a new FlowSpec capability which need to be advertised in an RI Opaque LSA. A new informational capability bit needs to be assigned for OSPF FlowSpec feature (FlowSpec Bit: TBD3).

This document defines a new BGP extended community attribute, i.e. import-policy (Type Code: TBD4), which is used to indicate whether importing the FlowSpec routes to the IGP protocol or not.

6. Security considerations

This extension to OSPF does not change the underlying security issues inherent in the existing OSPF. Implementations must assure that malformed TLV and Sub-TLV permutations do not result in errors which cause hard OSPF failures.

7. Acknowledgement

TBD.

8. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, February 2006.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.

[RFC4970] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, July 2007.

[RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, August 2009.

Authors' Addresses

Qiandeng Liang

Huawei
101 Software Avenue, Yuhuatai District
Nanjing, 210012
China

Email: liuweihang@huawei.com

Jianjie You
Huawei
101 Software Avenue, Yuhuatai District
Nanjing, 210012
China

Email: youjianjie@huawei.com