Ospf Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 24, 2015

Q. Liang
J. You
Huawei
September 20, 2014

OSPF Extensions for Flow Specification
draft-liang-ospf-flowspec-extensions-00

Abstract

This document defines a new OSPF flow specification (FlowSpec) Opaque Link State Advertisement (LSA) encoding format that can be used to distribute traffic flow specifications.

Additionally, this document proposes to assign label for the FlowSpec route in order to generate one or more LSPs, which would improve the route lookup performance in the data plane.

For the network not deploying BGP or the internal nodes in an AS, it is expected to extend IGP (Interior Gateway Protocol) to distribute FlowSpec routes. One application is to mitigate denial-of-service attacks as close to the attacker as possible, the other is to generate LSPs for the FlowSpec routes in the network.

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 24, 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 can be used 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 attacking traffic more effectively, i.e. block the attracking traffic as close to the attacker as possible, it is better to distribute the BGP FlowSpec rules to the customer network.

This document defines a new OSPF flow specification (FlowSpec) Opaque Link State Advertisement (LSA) [RFC5250] encoding format that can be used to distribute traffic flow specifications to the edge routers in the customer network. This mechanism can be used to mitigate denial-

of-service attacks, i.e. the execution point of FlowSpec rules would be much closer to the attacker.

Additionally, this document proposes to assign label for the FlowSpec route in order to generate one or more LSPs, which would improve the route lookup performance in the data plane. In this way, only the ingress LSR needs to identify a particular traffic flow based on the FlowSpec rules, e.g. ACL (Access Control List), and steer the packet to an LSP. Other LSRs of the LSP just need to forward the packets according to the label carried in the packets.

For the network not deploying BGP or the internal nodes in an AS, it is expected to extend IGP to distribute FlowSpec routes. One application is to mitigate denial-of-service attacks from the source of the attack. The other is to generate LSPs for the FlowSpec route in the network, thus an exclusive path could be set up for the traffic flow identified by this FlowSpec in order for the refined traffic tracking and statistics.

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. OSPF Extensions for FlowSpec Rules

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 1:

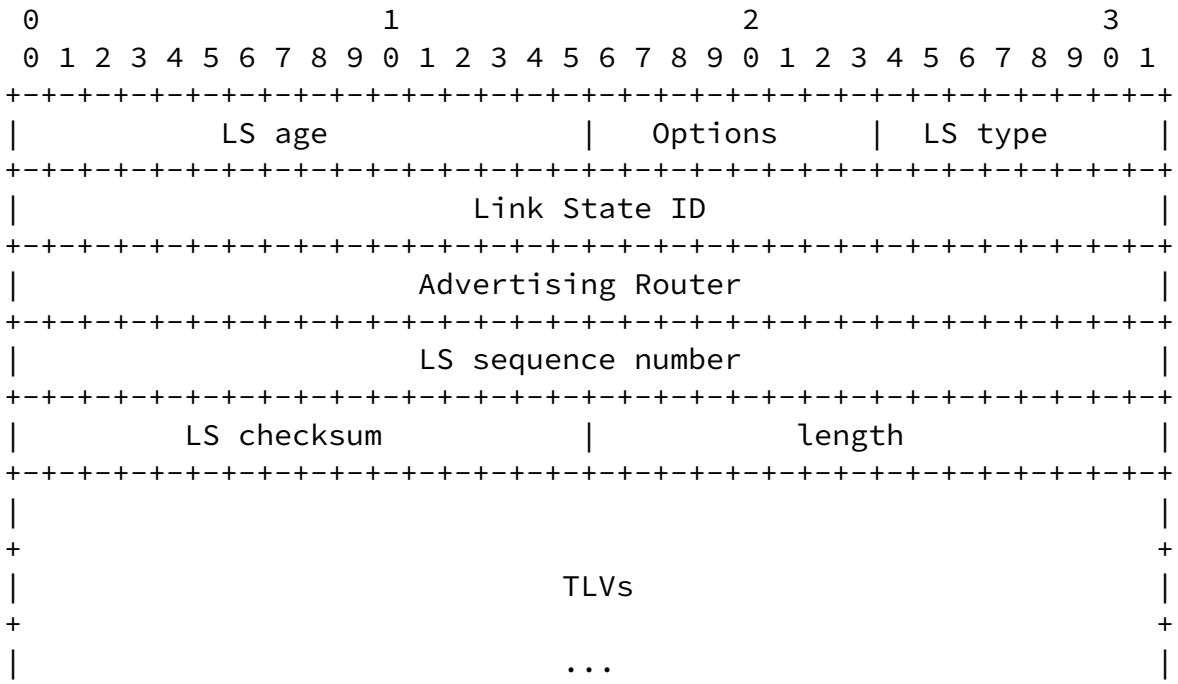

Figure 1: FlowSpec Opaque LSA

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 2:

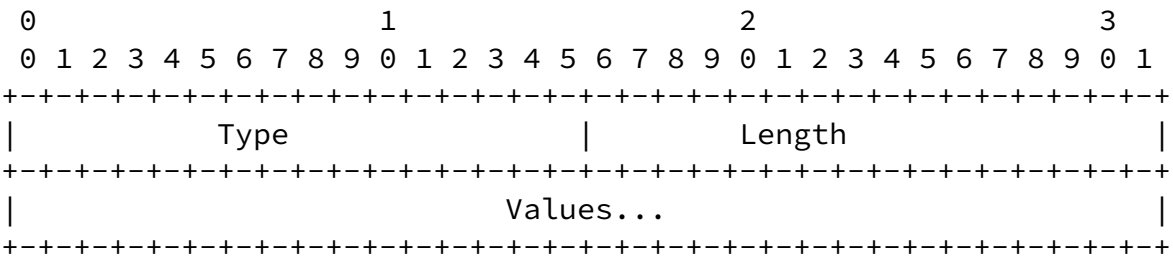

Figure 2: TLV Format

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.

3.1. OSPF FlowSpec Filter TLV

OSPF FlowSpec Filter TLV is carried in the FlowSpec Opaque LSA. It is defined below in Figure 3. There's one and only one OSPF FlowSpec Filter TLV included in a FlowSpec Opaque LSA.

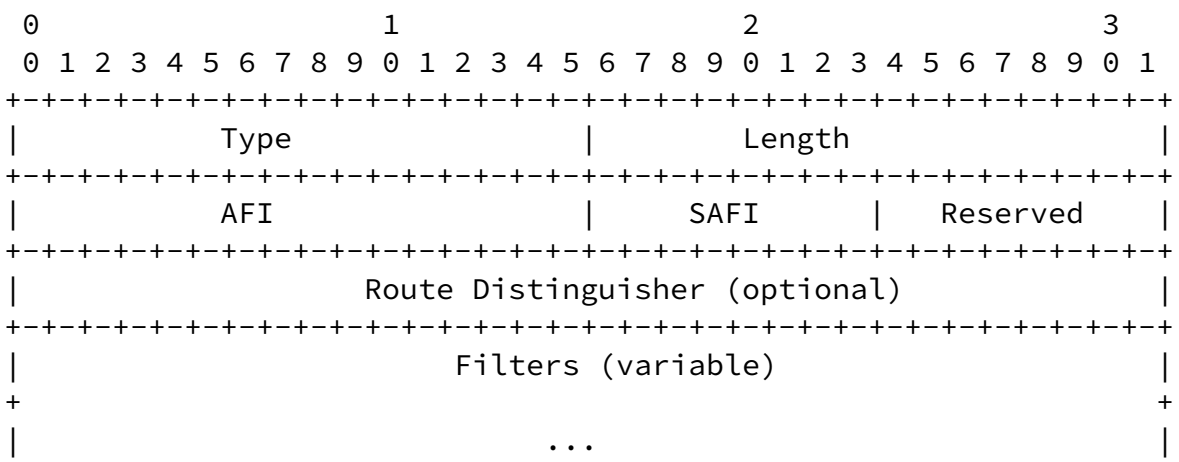

Figure 3: OSPF FlowSpec Filter TLV

Type: the TLV type (Type Code: TBD2)

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

AFI/SAFI: a particular application is identified by a specific (Address Family Identifier, Subsequent Address Family Identifier (AFI, SAFI)) pair [RFC4760] and corresponds to a distinct set of RIBs. Those RIBs should be treated independently from each other in order to assure non-interference between distinct applications.

Route Distinguisher: when AFI/SAFI (e.g. AFI=1, SAFI=134) indicates VPN service, this field is available, otherwise, it is not present.

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

3.2. OSPF FlowSpec Action TLV

This document defines a set of actions associated with a particular

FlowSpec rule. There's one or more OSPF FlowSpec Action TLVs with different types included in a FlowSpec Opaque LSA. The following OSPF FlowSpec Action TLVs (Table 1) 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

Two new OSPF FlowSpec Action TLVs are defined below in Table 2:

Table 2: Extended FlowSpec Actions

type	FlowSpec Action	encoding
TBD3	traffic-nexthop	IPv4 or IPv6
TBD4	traffic-label	3-byte Lable and EXP.

Traffic-nexthop: a next hop IP address instructed by FlowSpec route information. The default value is the local IP address of

the current router, i.e. the one sending the OSPF message. It is 4 bytes for IPv4, and 16 bytes for IPv6.

Traffic-label: a label assigned by OSPF to a FlowSpec route, in order to generate one or more LSPs. It is 3 bytes for label and EXP.

3.3. 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:

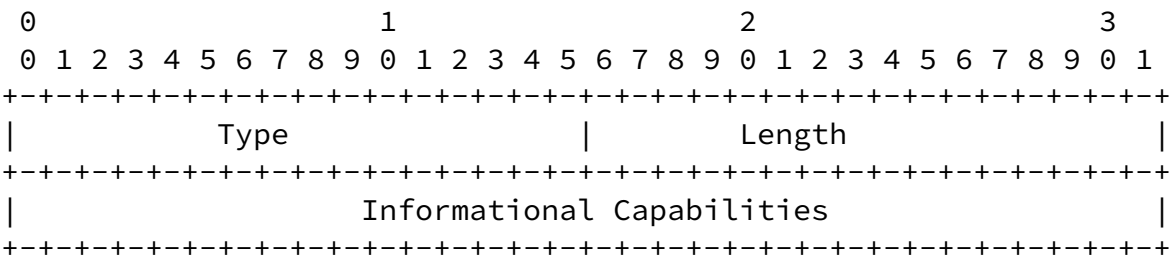

Figure 4: OSPF RI Capabilities TLV

The following informational capability bits are assigned:

Bit	Capabilities
6 (TBD5)	OSPF FlowSpec
7-31	Unassigned (Standards Action)

3.4. 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 action 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: TBD6) is defined as follows:

type	extended community	encoding
TBD	import-policy	IGP target

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

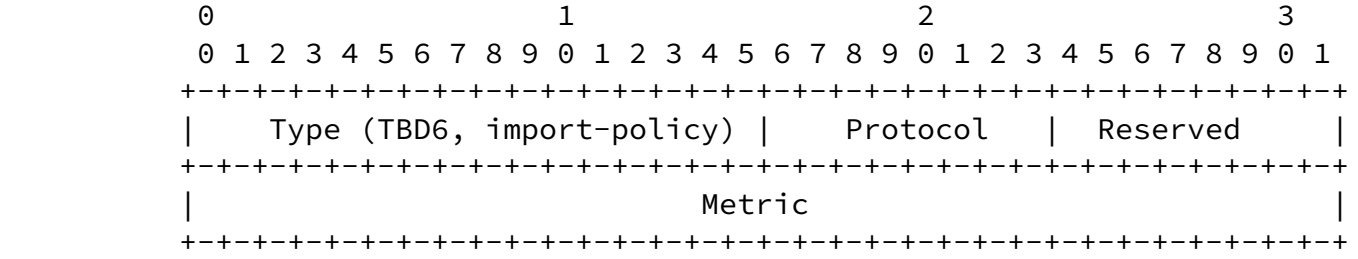

Figure 5: Import-policy Extended Community

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.

4. 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 Filter TLV (Type Code: TBD2), which is used to describe the filters.

This document defines two OSPF FlowSpec Action TLVs, i.e. traffic-nexthop (Type Code: TBD3) and traffic-label (Type Code: TBD4), which are 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: TBD5).

This document defines a new BGP extended community attribute, i.e. import-policy (Type Code: TBD6), which is used to indicate whether importing the FlowSpec routes to the IGP protocol or not.

5. 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.

6. Acknowledgement

TBD.

7. 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