

Ospf Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 16, 2016

Q. Liang
J. You
N. Wu
Huawei
P. Fan
Independent
K. Patel
A. Lindem
Cisco Systems
April 14, 2016

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

Abstract

Dissemination of the Traffic flow information was first introduced in the BGP protocol [[RFC5575](#)]. FlowSpec routes are used to distribute traffic filtering rules that are used to filter Denial-of-Service (DoS) attacks. For the networks that only deploy an IGP (Interior Gateway Protocol) (e.g., OSPF), it is required that the IGP is extended to distribute Flow Specification or FlowSpec routes.

This document discusses the use cases for distributing flow specification (FlowSpec) routes using OSPF. Furthermore, this document defines a OSPF FlowSpec Opaque Link State Advertisement (LSA) encoding format that can be used to distribute FlowSpec routes, its validation procedures for imposing the filtering information on the routers, and a capability to indicate the support of FlowSpec functionality.

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-Draft

OSPF FlowSpec

April 2016

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 October 16, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://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.	Terminology	3
3.	Use Cases for OSPF based FlowSpec Distribution	4
3.1.	OSPF Campus Network	4
3.2.	BGP/MPLS VPN	4
3.2.1.	Traffic Analyzer Deployed in Provider Network	5
3.2.2.	Traffic Analyzer Deployed in Customer Network	6
3.2.3.	Policy Configuration	6
4.	OSPF Extensions for FlowSpec Rules	7
4.1.	FlowSpec LSA	7
4.1.1.	OSPFv2 FlowSpec Opaque LSA	7
4.1.2.	OSPFv3 FlowSpec LSA	9
4.2.	OSPF FlowSpec Filters TLV	10
4.2.1.	Interface-Set TLV	11
4.2.2.	Order of Traffic Filtering Rules	12
4.2.3.	Validation Procedure	12
4.3.	OSPF FlowSpec Action TLV	13
4.3.1.	Traffic-rate	14
4.3.2.	Traffic-action	14

4.3.3.	Traffic-marking	14
4.3.4.	Redirect-to-IP	15
4.4.	Capability Advertisement	16
5.	Redistribution of FlowSpec Routes	16
6.	IANA Considerations	16

7.	Security considerations	17
8.	Acknowledgement	17
9.	References	17
9.1.	Normative References	17
9.2.	Informative References	18
	Authors' Addresses	19

[1.](#) Introduction

[RFC5575] defines Border Gateway Protocol protocol extensions that can be used to distribute traffic flow specifications. One application of this 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 networks deploying only an IGP (e.g., OSPF), it is expected to extend the IGP (OSPF in this document) to distribute FlowSpec routes. This document discusses the use cases for distributing FlowSpec routes using OSPF. Furthermore, 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, its validation procedures for imposing the filtering information on the routers, and a capability to indicate the support of Flowspec functionality.

The semantic content of the FlowSpec extensions defined in this document are identical to the corresponding extensions to BGP ([RFC5575] and [I-D.ietf-idr-flow-spec-v6]). In order to avoid repetition, this document only concentrates on those parts of specification where OSPF is different from BGP. The OSPF flowspec extensions defined in this document 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 networks deploying only an IGP (e.g., OSPF), it is expected to extend the IGP (OSPF in this document) to distribute FlowSpec routes, because when the FlowSpec routes are installed in the customer network, they are 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 route distribution.

[3.1.](#) OSPF Campus Network

For networks 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 deployed with a router, then the FlowSpec routes from the traffic analyzer need to be distributed to the other routers in this domain using OSPF.

```
+-----+
|Traffic |
+---+Analyzer|
|   +-----+
|
|FlowSpec
|
```

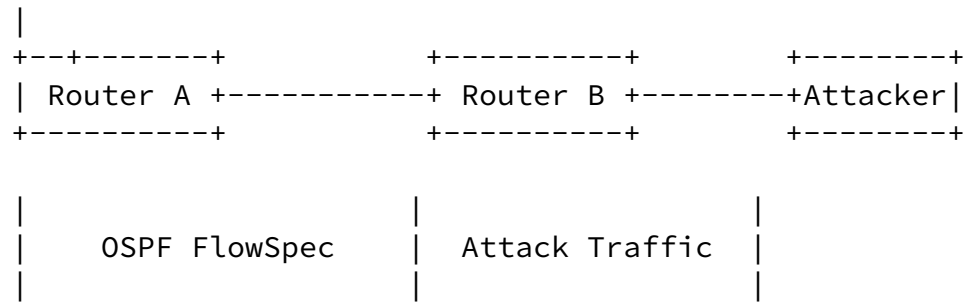


Figure 3: OSPF Campus Network

3.2. 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, or OSPF) is used between the PE (Provider Edge) and CE (Customer Edge) in many deployments. In order to distribute the FlowSpec routes to the customer network, the IGP needs to support 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 are discussed below.

3.2.1. 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. FlowSpec routes with a Route Distinguisher (RD) in the Network Layer Reachability information (NRLI) corresponding to VPN1 are distributed from the traffic analyzer to the PE1 to which the traffic analyzer is attached. If the traffic analyzer is also a BGP speaker, it can distribute the FlowSpec routes using BGP [RFC5575]. Then the PE1 distributes the FlowSpec routes further to the PE2. Finally, the FlowSpec routes need to be distributed from PE2 to the CE2 using OSPF, i.e., to the customer network VPN1. As an attacker is more likely in the customer network, FlowSpec routes installed directly on CE2 could mitigate the impact of DoS attacks better.

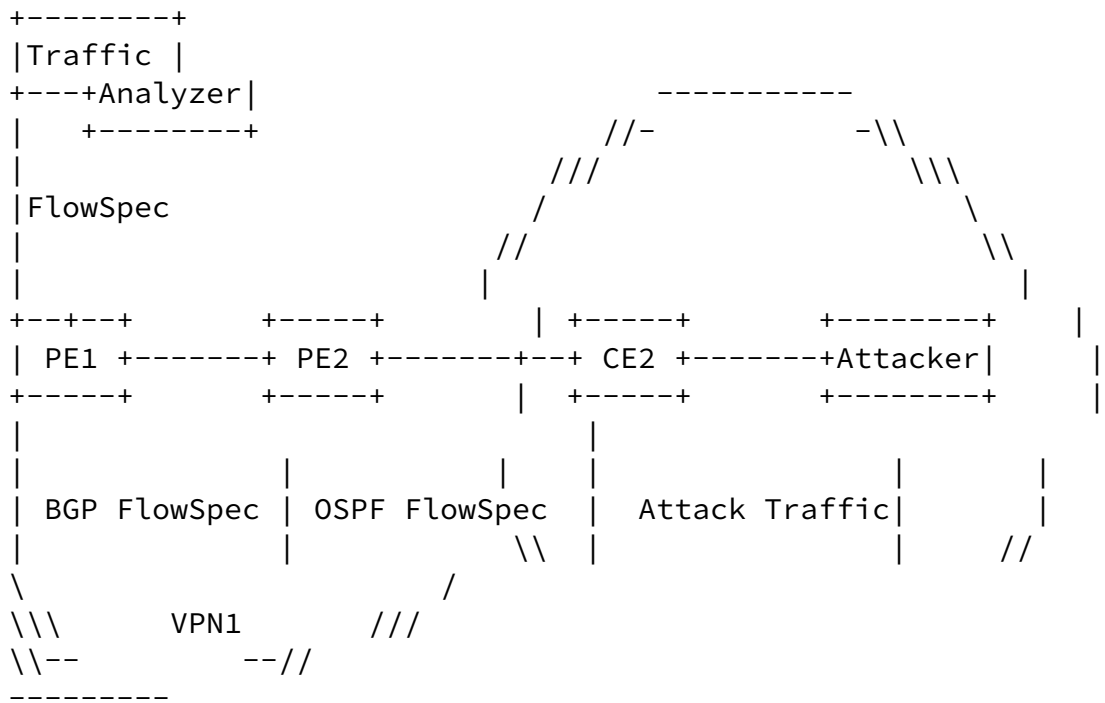


Figure 1: Traffic Analyzer deployed in Provider Network

[3.2.2.](#) 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 to prevent associated DoS attacks. Then the FlowSpec routes would be distributed from the traffic analyzer to the CE1 using OSPF or another policy protocol (e.g., RESTful API over HTTP). Furthermore, the FlowSpec routes need to be distributed throughout the provider network via PE1/PE2 to 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.

+-----+

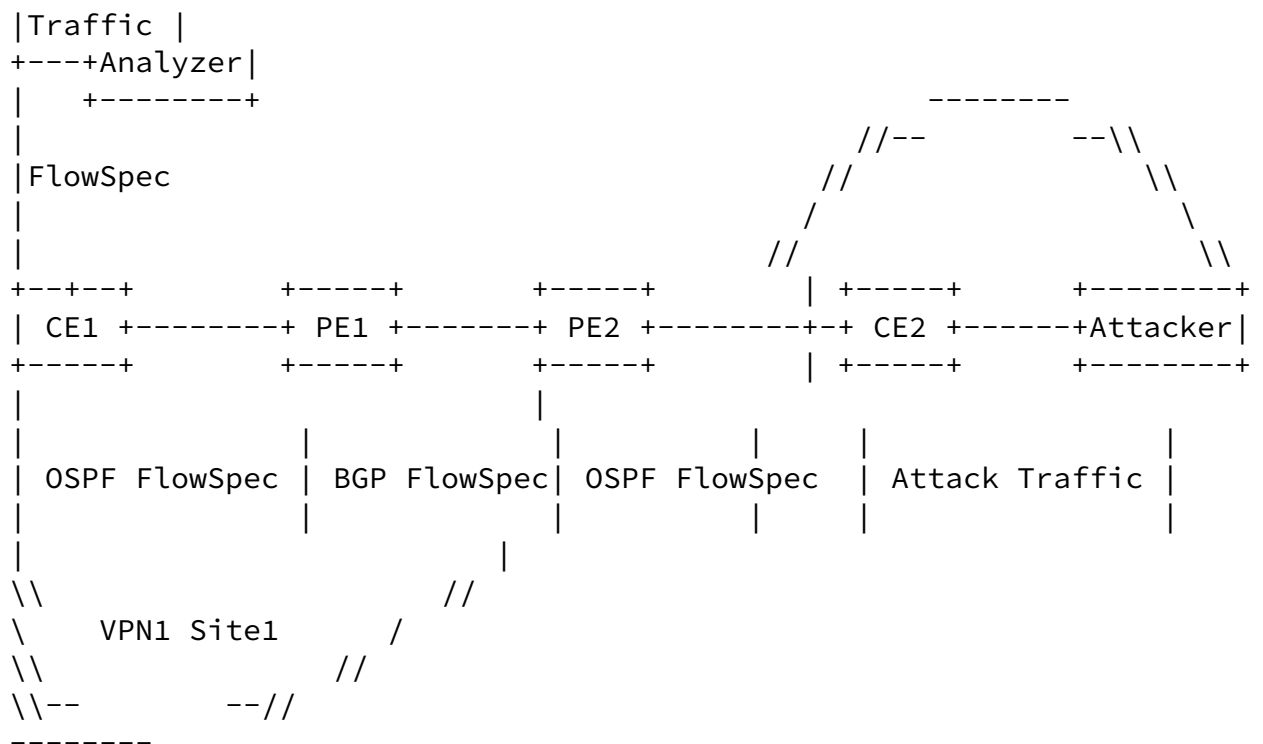


Figure 2: Traffic Analyzer deployed in Customer Network

3.2.3. Policy Configuration

The CE or PE could deploy local filtering policies to filter OSPF FlowSpec rules, for example, deploying a filtering policy to filter the incoming OSPF FlowSpec rules in order to prevent illegal or invalid FlowSpec rules from being applied.

The PE should configure FlowSpec importing policies to control importing action between the BGP IP/VPN FlowSpec RIB and the OSPF Instance FlowSpec RIB. Otherwise, the PE couldn't transform a BGP

IP/VPN FlowSpec rule to an OSPF FlowSpec rule or transform an OSPF FlowSpec rule to a BGP IP/VPN FlowSpec rule.

4. OSPF Extensions for FlowSpec Rules

4.1. FlowSpec LSA

4.1.1. OSPFv2 FlowSpec Opaque LSA

areas, which may aggravate the burden of devices in that area.

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-octet 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.

If FlowSpec Opaque LSA is Type-11 Opaque LSA, it is not flooded into Stub and NSSA areas. As the traffic accessing a network segment outside Stub and NSSA areas would be aggregated to the ABR, FlowSpec rules could be applied on the ABR instead of disseminating them into Stub and NSSA areas.

4.1.2. OSPFv3 FlowSpec LSA

This document defines a new OSPFv3 flow specification LSA encoding format that can be used to distribute traffic flow specifications. This new OSPFv3 FlowSpec LSA is extended based on [\[RFC5340\]](#).

The OSPFv3 FlowSpec LSA is defined below in Figure 6:

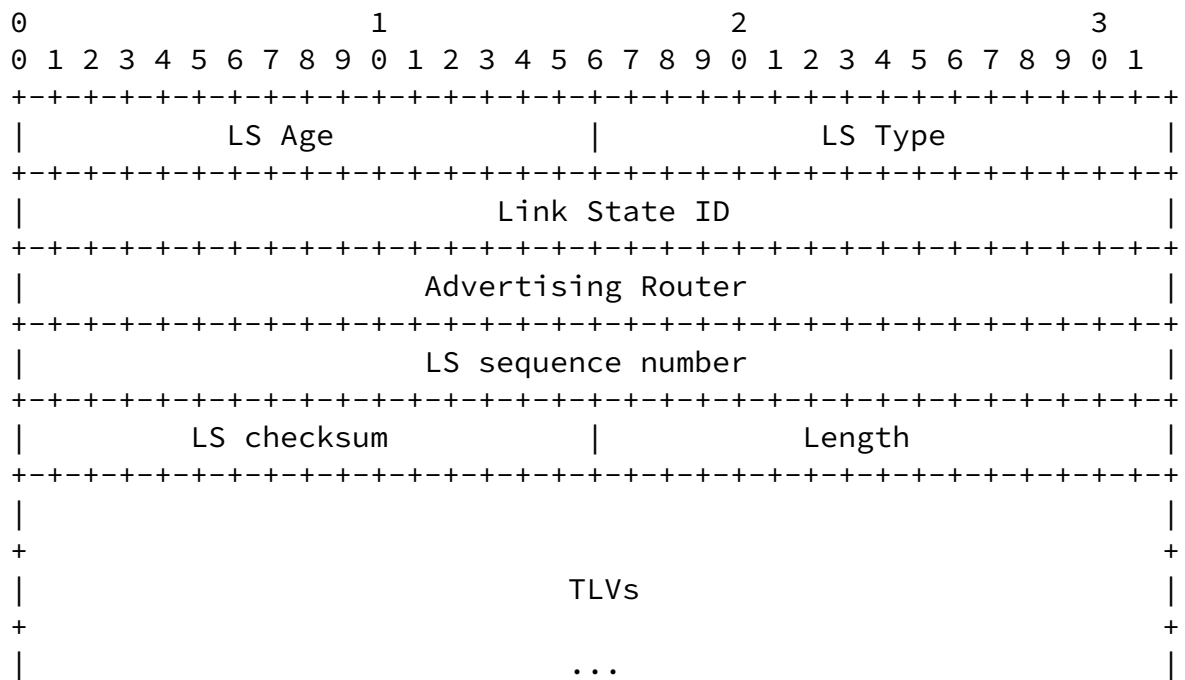


Figure 6: OSPFv3 FlowSpec LSA

LS age: the same as defined in [\[RFC5340\]](#).

LS type: the same as defined in [\[RFC5340\]](#). The format of the LS type is as follows:




Figure 7: LSA Type

In this document, the U bit should be set indicating that the OSPFv3 FlowSpec LSA should be flooded even if it is not understood. For the area scope, the S1 bit should be set and the S2 should be clear. For the AS scope, the S1 bit should be clear

and the S2 bit should be set. A new LSA Function Code (TBD2) needs to be defined for OSPFv3 FlowSpec LSA. To facilitate inter-

area reachability validation, any OSPFv3 router originating AS scoped LSAs is considered an AS Boundary Router (ASBR).

Link State ID: the same as defined in [RFC5340].

Advertising Router: the same as defined in [RFC5340].

LS sequence number: the same as defined in [RFC5340].

LS checksum: the same as defined in [RFC5340].

Length: the same as defined in [RFC5340].

TLVs: one or more TLVs MAY be included in a OSPFv3 FlowSpec LSA to carry FlowSpec information.

4.2. OSPF FlowSpec Filters TLV

The FlowSpec Opaque LSA carries one or more FlowSpec Filters TLVs and corresponding FlowSpec Action TLVs. The OSPF FlowSpec Filters TLV is defined below in Figure 8.




Figure 8: OSPF FlowSpec Filters TLV

Type: the TLV type (Type Code: TBD3)

Length: the size of the value field in octets

Flags: One octet Field identifying Flags.

```

      0 1 2 3 4 5 6 7
    +---+---+---+---+---+---+
    | Reserved      |S|
    +---+---+---+---+---+---+
```

The least significant bit S is defined as a strict Filter check bit. If set, Strict Validation rules outlined in the validation [Section 4.2.2](#) need to be enforced.

Filters: the same as "flow-spec NLRI value" defined in [[RFC5575](#)] and [[I-D.ietf-idr-flow-spec-v6](#)].

Table 1: OSPF Supported FlowSpec Filters

Type	Description	RFC/ WG draft
1	Destination IPv4 Prefix Destination IPv6 Prefix	RFC5575 I-D.ietf-idr-flow-spec-v6
2	Source IPv4 Prefix Source IPv6 Prefix	RFC5575 I-D.ietf-idr-flow-spec-v6
3	IP Protocol Next Header	RFC5575 I-D.ietf-idr-flow-spec-v6
4	Port	RFC5575
5	Destination port	RFC5575
6	Source port	RFC5575
7	ICMP type	RFC5575
8	ICMP code	RFC5575
9	TCP flags	RFC5575

10	Packet length	RFC5575	
+-----+		+-----+	
11	DSCP	RFC5575	
+-----+		+-----+	
12	Fragment	RFC5575	
+-----+		+-----+	
13	Flow Label	I-D.ietf-idr-flow-spec-v6	
+-----+		+-----+	
14	Interface-Set	Described Below	
+-----+		+-----+	

[4.2.1.](#) Interface-Set TLV

The Interface-Set TLV is used to limit the FlowSpec rules to a set of interfaces configured locally with the specified Group ID. The Interface-Set TLV was inspired by

Liang, et al.

Expires October 16, 2016

[Page 11]

Internet-Draft

OSPF FlowSpec

April 2016

[I-D.litkowski-idr-flowspec-interfaceset] and uses similar encodings. The Autonomous System (AS) number is not required since OSPF usage is within a single AS.

The Interface-Set TLV is encoded as:

0	1	2	3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1			
+-----+		+-----+	
TBD, 14 Suggested	4		
+-----+		+-----+	
0 I Flags	Group ID		
+-----+		+-----+	

0 : if set, the flow specification rule MUST be applied in outbound direction to the interface set referenced by the specified Group ID.

I : if set, the flow specification rule MUST be applied in input direction to the interface set referenced by the specified Group ID

Both flags can be set at the same time in the interface-set extended community leading to flow rule to be applied in both directions. An interface-set TLV with both flags set to zero MUST be treated as an error and as consequence, the FlowSpec update MUST be ignore and an error should be logged.

The Group Identifier is coded as a 16-bit number (values goes from 0 to 65535).

Multiple instances of the interface-set community may be present in a Flow-Spec rule. This may appear if the flow rule need to be applied to multiple set of interfaces.

[4.2.2.](#) Order of Traffic Filtering Rules

With traffic filtering rules, more than one rule may match a particular traffic flow. The order of applying the traffic filter rules is the same as described in [Section 5.1 of \[RFC5575\]](#) and in Section 3.1 of [\[I-D.ietf-idr-flow-spec-v6\]](#).

[4.2.3.](#) Validation Procedure

[RFC5575] defines a validation procedure for BGP FlowSpec rules, and [\[I-D.ietf-idr-bgp-flowspec-oid\]](#) describes a modification to the validation procedure defined in [\[RFC5575\]](#) for the dissemination of BGP flow specifications. The OSPF FlowSpec should support similar features to mitigate the unnecessary application of traffic filter

rules. The OSPF FlowSpec validation procedure is described as follows.

When a router receives a FlowSpec rule including a destination prefix filter from its neighbor router, it should consider the prefix filter as a valid filter unless the S bit in the flags field of Filter TLV is set. If the S bit is set, then the FlowSpec rule is considered valid if and only if:

The originator of the FlowSpec rule matches the originator of the best-match unicast route for the destination prefix embedded in the FlowSpec.

The former rule allows any centralized controller to originate the prefix filter and advertise it within a given OSPF network. The latter rule, also known as a Strict Validation rule, allows strict checking and enforces that the originator of the FlowSpec filter is also the originator of the destination prefix.

When multiple equal-cost paths exist in the routing table entry, each path could end up having a separate set of FlowSpec rules.

When a router receives a FlowSpec rule not including a destination prefix filter from its neighbor router, the validation procedure described above is not applicable.

The FlowSpec filter validation state is used by a speaker when the filter is considered for an installation in its FIB. An OSPF speaker MUST flood OSPF FlowSpec LSA as per the rules defined in [RFC2328] regardless of the validation state of the prefix filters.

4.3. OSPF FlowSpec Action TLV

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

Redirect: IPv4 or IPv6 address. This IP address may correspond to a tunnel, i.e., the redirect allows the traffic to be redirected to a directly attached next-hop or a next-hop requiring a route lookup.

Table 2: Traffic Filtering Actions in [RFC5575], etc.

type	FlowSpec Action	RFC/WG draft
0x8006	traffic-rate	RFC5575
0x8007	traffic-action	RFC5575
0x8108	redirect-to-IPv4	I-D.ietf-idr-flowspec-redirect-rt-bis
0x800b	redirect-to-IPv6	I-D.ietf-idr-flow-spec-v6

0x8009	traffic-marking	RFC5575
--------	-----------------	-------------------------

[4.3.1.](#) Traffic-rate

Traffic-rate TLV is encoded as:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
TBD5, 0x8006 suggested										4																													
Traffic-rate																																							

Traffic-rate: the same as defined in [[RFC5575](#)].

[4.3.2.](#) Traffic-action

Traffic-action TLV is encoded as:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
TBD6, 0x8007 suggested										2																													
Reserved										S T																													

S flag and T flag: the same as defined in [[RFC5575](#)].

[4.3.3.](#) Traffic-marking

Traffic-marking TLV is encoded as:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
TBD7, 0x8009 suggested										2																													


```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Reserved   | DSCP Value|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

DSCP value: the same as defined in [[RFC5575](#)].

[4.3.4.](#) Redirect-to-IP

Redirect-to-IPv4 is encoded as:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   TBD8, 0x8108 suggested   |           6           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv4 Address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Reserved                   |C|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Redirect to IPv6 TLV is encoded as (Only for OSPFv3):

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   TBD9, 0x800b suggested   |          18          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv6 Address                               |
|                               |
|                               |
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Reserved                   |C|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

IPv4/6 Address: the redirection target address.

'C' (or copy) bit: when the 'C' bit is set, the redirection applies to copies of the matching packets and not to the original traffic stream [[I-D.ietf-idr-flowspec-redirect-ip](#)].

[4.4.](#) Capability Advertisement

This document defines a capability bit for OSPF Router-Information LSA [[RFC7770](#)] as FlowSpec Capability Advertisement bit. When set, the OSPF router indicates its ability to support the FlowSpec functionality. The FlowSpec Capability Advertisement bit has a value to be assigned by IANA from OSPF Router Functional Capability Bits Registry [I-D.ietf-ospf-rfc4970bis].

[5.](#) Redistribution of FlowSpec Routes

In certain scenarios, FlowSpec routes MAY get redistributed from one protocol domain to another; specifically from BGP to OSPF and vice-versa. When redistributed from BGP, the OSPF speaker SHOULD generate an Opaque LSA for the redistributed routes and announce it within an OSPF domain. An implementation MAY provide an option for an OSPF speaker to announce a redistributed FlowSpec route within a OSPF domain regardless of being installed in its local FIB. An implementation MAY impose an upper bound on number of FlowSpec routes that an OSPF router MAY advertise.

[6.](#) IANA Considerations

This document defines a new OSPFv2 Opaque LSA, i.e., OSPFv2 FlowSpec Opaque LSA (Type Code: TBD1), which is used to distribute traffic flow specifications.

This document defines a new OSPFv3 LSA, i.e., OSPFv3 FlowSpec LSA (LSA Function Code: TBD2), which is used to distribute traffic flow specifications.

This document defines OSPF FlowSpec Filters TLV (Type Code: TBD3), 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: TBD4).

This document defines a new Router LSA bit known as a FlowSpec Capability Advertisement bit. This document requests IANA to assign a bit code type for FlowSpec Capability Advertisement bit from the OSPF Router Functional Capability Bits registry.

- Type 1 - Destination IPv4/IPv6 Prefix
- Type 2 - Source IPv4/IPv6 Prefix
- Type 3 - IP Protocol/Next Header
- Type 4 - Port
- Type 5 - Destination port
- Type 6 - Source port
- Type 7 - ICMP type
- Type 8 - ICMP code
- Type 9 - TCP flags
- Type 10 - Packet length
- Type 11 - DSCP
- Type 12 - Fragment
- Type 13 - Flow Label
- Type 14 - Interface-Set

This document defines a group of FlowSpec actions. The following TLV types need to be assigned:

- Type 0x8006(TBD5) - traffic-rate
- Type 0x8007(TBD6) - traffic-action
- Type 0x8009(TBD7) - traffic-marking
- Type 0x8108(TBD8) - redirect to IPv4
- Type 0x800b(TBD9) - redirect to IPv6

[7.](#) 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.

[8.](#) Acknowledgement

The authors would also like to thank Burjiz Pithawala, Rashmi Shrivastava and Mike Dubrovsky for their contribution to the original version of the document.

[9.](#) References

[9.1.](#) 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, <<http://www.rfc-editor.org/info/rfc2119>>.

Liang, et al.

Expires October 16, 2016

[Page 17]

Internet-Draft

OSPF FlowSpec

April 2016

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), DOI 10.17487/RFC5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), DOI 10.17487/RFC5340, July 2008, <<http://www.rfc-editor.org/info/rfc5340>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](#), DOI 10.17487/RFC5575, August 2009, <<http://www.rfc-editor.org/info/rfc5575>>.

[9.2](#). Informative References

- [I-D.ietf-idr-bgp-flowspec-oid]
Uttaro, J., Filsfils, C., Smith, D., Alcaide, J., and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications", [draft-ietf-idr-bgp-flowspec-oid-03](#) (work in progress), March 2016.
- [I-D.ietf-idr-flow-spec-v6]
McPherson, D., Raszuk, R., Pithawala, B., Andy, A., and S. Hares, "Dissemination of Flow Specification Rules for IPv6", [draft-ietf-idr-flow-spec-v6-07](#) (work in progress), March 2016.
- [I-D.ietf-idr-flowspec-redirect-ip]
Uttaro, J., Haas, J., Texier, M., Andy, A., Ray, S.,

Simpson, A., and W. Henderickx, "BGP Flow-Spec Redirect to IP Action", [draft-ietf-idr-flowspec-redirect-ip-02](#) (work in progress), February 2015.

[I-D.litkowski-idr-flowspec-interfaceset]

Litkowski, S., Simpson, A., Patel, K., and J. Haas, "Applying BGP flowspec rules on a specific interface set", [draft-litkowski-idr-flowspec-interfaceset-03](#) (work in progress), December 2015.

[RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 7770](#), DOI 10.17487/RFC7770, February 2016, <<http://www.rfc-editor.org/info/rfc7770>>.

Liang, et al.

Expires October 16, 2016

[Page 18]

Internet-Draft

OSPF FlowSpec

April 2016

Authors' Addresses

Qiandeng Liang
Huawei
101 Software Avenue, Yuhuatai District
Nanjing, 210012
China

Email: liangqiandeng@huawei.com

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

Email: youjianjie@huawei.com

Nan Wu
Huawei

Email: eric.wu@huawei.com

Peng Fan
Independent

Email: peng.fan@139.com

Keyur Patel
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: keyupate@cisco.com

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27519
USA

Email: acee@cisco.com