

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 19, 2013

R. Shrivastava
M. Dubrovsky
K. Patel
B. Pithawala
Cisco Systems
April 17, 2013

Flow Specification Extensions to OSPF Protocol
draft-shrivastava-ospf-flow-spec-01

Abstract

This document describes the extensions to the OSPF protocol to distribute the traffic flow specification rules and associated actions. The notion and principles of operation of the mechanism were initially introduced into the BGP protocol. The IPv4 part of specification was published in [RFC5575](#) and later the specification was enhanced for IPv6 via [draft-raszuk-idr-flow-spec-v6](#). The mechanism allows the routing protocol to distribute information about the traffic flows and associated actions such as packet filtering with it.

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 [RFC 2119](#) [[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 October 19, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<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.	The Flow Specification Operation Description	4
2.1.	The Flow Specification Origins	4
2.2.	Flow-specification-LSA scope and re-origination	4
2.3.	IPv4 and IPv6 rules	4
3.	OSPFv2 Flow-specification-LSA	5
3.1.	Link-State ID format	5
4.	OSPFv3 Flow-specification-LSA	5
5.	Flow-specification-LSA payload	6
6.	IPv4 Flow Specification	6
6.1.	Destination Prefix TLV	7
6.2.	Source Prefix TLV	7
6.3.	IP Protocol	7
6.4.	Port	7
6.5.	Destination port	7
6.6.	Source port	8
6.7.	ICMP Type	8
6.8.	ICMP code	8
6.9.	TCP flags	8
6.10.	Packet length	8
6.11.	DSCP (Diffserv Code Point)	8
6.12.	Fragment	8
7.	IPv6 Flow Specification	9
7.1.	Destination IPv6 Prefix TLV	9
7.2.	Source IPv6 Prefix TLV	10
7.3.	Traffic Class	10
7.4.	Flow Label	10
8.	IPv4 Traffic Filtering Actions	10

8.1.	Traffic-rate	11
8.2.	Traffic-action	11
8.3.	Redirect	11
8.4.	Traffic-marking	11
9.	IPv6 Traffic Filtering Actions	12

9.1.	Traffic-marking	12
10.	Order of Traffic Filtering Rules	12
11.	Validation Procedure	12
12.	Security Considerations	13
13.	IANA Considerations	13
14.	Acknowledgements	15
15.	Normative References	15
	Authors' Addresses	16

[1.](#) Introduction

This document describes a method of adding capability to distribute traffic flow specification rules into OSPF.

The semantic content of the extensions is identical to the corresponding extensions to BGP ([[BGP-FLOWSPEC](#)] and [[draft-raszuk-idr-flow-spec-v6](#)]). In order to avoid repetition, this document only concentrates on those parts of specification where OSPF is different from BGP.

Each flow specification rules is encapsulated into the flow-specification-LSA and then flooded along with the prefix reachability information across an area or the entire OSPF routing domain. When a router receives the LSA from its neighbor, it decodes the data, validates the information and applies the filtering rules as defined.

Some OSPF routers can be configured to originate flow specification rules with associated actions. Other routers can be configured to apply these rules from specific sets of OSPF Router IDs.

The use cases for the traffic policy mechanism are different from the BGP. Let's look at some examples.

Example : Installation of a new voice gateway in the enterprise network.

The router, which is directly connected to the new voice gateway, originates flow specification rules that describe voice traffic and set QoS marking action. Routers on the other side of the WAN links are configured to act on this information.

Another example: temporarily contractors.

The contractor company, which is connected to the enterprise network via the WAN link, should have limited access to certain servers. The routers directly connected to the corporate data center can originate rules describing services accessibility for contractors. The edge routers with WAN links to the contractor's equipment install those rules to block the undesirable traffic.

Last example: Service provider MPLS VPN hub-and-spoke solution for the enterprise customer.

The spokes are connected to the MPLS VPN backbone with T1 lines and the hub with OC3. OSPF CE spoke router generates flow spec with rate-limit restrictions for spoke aggregated ip prefix. The flow specification is then transmitted over the MPLS VPN core. OSPF CE hub router uses the rule to limit the traffic towards the spoke.

This document does not address optimizations in storage of flow-specification-LSAs in the intermediate routers or auto-discovery of flow specification capable routers.

[2.](#) The Flow Specification Operation Description

[2.1.](#) The Flow Specification Origins

Flow specification rules can be either originated by the same router that originates the corresponding destination prefix. For example, in case of an OSPFv2 network-LSA, the Designated Router will also advertise the corresponding flow specification rules that apply to the prefix advertised in the network-LSA.

Each flow specification rule is encapsulated into a separate flow-specification-LSA that is described below in [Section 3](#) and [Section 4](#).

The receiving routers first validate and then install flow specification rules into the flow routing table.

[2.2.](#) Flow-specification-LSA scope and re-origination

Because of the validation restrictions described in [Section 11](#) the flow-specification-LSA MUST have the same scope as the scope of the corresponding LSA that carries the destination prefix information of the flow. When the LSA describing the destination prefix is translated by the ABR router, this router SHOULD also re-originate the corresponding flow-specification-LSAs with the same scope as the LSA carrying the translated prefix.

[2.3.](#) IPv4 and IPv6 rules

Shrivastava, et al. Expires October 19, 2013 [Page 4]

Internet-Draft OSPF Flow Specification April 2013

Currently the OSPFv2 specification [[OSPFV2](#)] supports only IPv4 addresses and can therefore only carry IPv4 flow specification rules. OSPFv3 [[OSPFV3-AF](#)] supports both IPv4 and IPv6 address families but in separate instances. Therefore a particular OSPFv3 instance can carry either IPv6 or IPv4 flow specification rules.


[3.](#) OSPFv2 Flow-specification-LSA

This extension makes use of the Opaque LSA [[OPAQUE](#)] to carry the flow specification rules. This proposal uses type 10 for area scope and type 11 for AS scope.

[3.1.](#) Link-State ID format

The link-state ID of the Opaque LSA is divided into an Opaque Type field (the first 8 bits) and an Opaque ID (the remaining 24 bits). The Flow-specification-LSA uses type TBD. The Opaque ID is an arbitrary value used to maintain multiple flow specification rules, which originate from a single router.


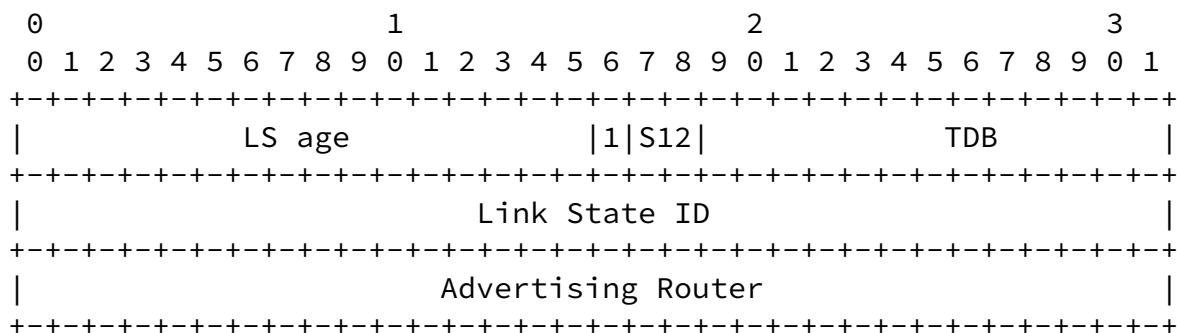
0	1	2	3
0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1



Flow-specification-LSA LSA ID format

4. OSPFv3 Flow-specification-LSA

The OSPFv3 Flow specification information will be carried in the LSA with function code TBD. The U bit will be set indicating that the OSPFv3 flow-specification-LSA should be flooded even if it is not understood. For the area scope flow-specification-LSA S1 bit should be set and S2 should be 0. For the AS scope, the S1 bit should be cleared and S2 bit should be set. Like the OSPFv2 Opaque ID field, the Link State ID is an arbitrary value used to maintain multiple flow specification rules originating from a single router.



OSPFv3 Flow-specification-LSA

5. Flow-specification-LSA payload

The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets as described in Section 2.3.2 of [\[MPLS-TE\]](#).

Each rule consists of one or more specification component types that identify the flow followed by optional flow specification filtering action types.

[6.](#) IPv4 Flow Specification

[BGP-FLOWSPEC] defines 12 flow specification component types:

Type	Description
-----	-----
1	Destination Prefix
2	Source Prefix
3	IP Protocol
4	Port
5	Destination Port
6	Source Port
7	ICMP type
8	ICMP code
9	TCP flags
10	Packet length
11	DSCP
12	Fragment

Each of the flow specification component types is defined in a separate TLV.

The flow specification component type is stored in the lower 8 bits of the 16-bit type field.

Shrivastava, et al. Expires October 19, 2013 [Page 6]

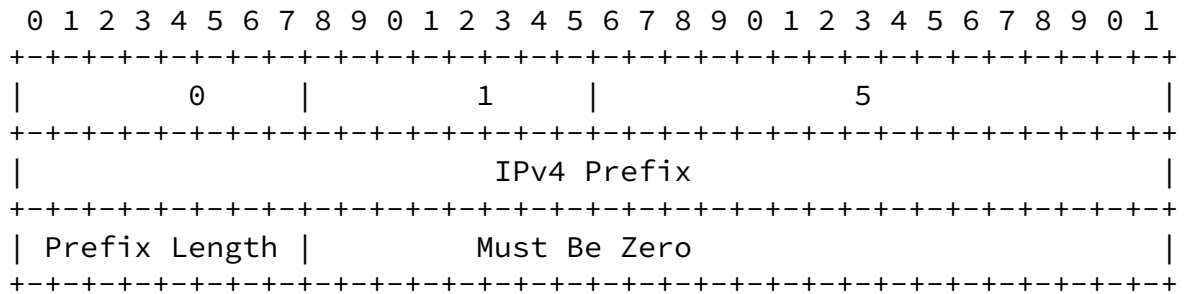
Internet-Draft OSPF Flow Specification April 2013

The format of each flow specification component type TLV is outlined below.

[6.1.](#) Destination Prefix TLV

Defines the destination prefix to match.


0 1 2 3



6.2. Source Prefix TLV

The encoding is similar to the destination prefix TLV. The type field is set to 2.

6.3. IP Protocol



The Value field has a variable length and the length is determined by the two bits len field in the operator byte.

6.4. Port

The same encoding as for component type IP Protocol. The TLV Type field is set to 4.

6.5. Destination port

The same encoding as for component type IP Protocol. The TLV type

6.6. Source port

The same encoding as for component type IP Protocol. The TLV type field is set to 6.


6.7. ICMP Type

The same encoding as for component type IP Protocol. The TLV Type field is set to 7.

6.8. ICMP code

The same encoding as for component type IP Protocol. The TLV Type field is set to 8.

6.9. TCP flags



The Bitmask field is of varying length and determined by the two bits len field in the operator byte.


6.10. Packet length

The same encoding as for component type IP Protocol. The TLV Type field is set to 10.

6.11. DSCP (Diffserv Code Point)

The same encoding as for component type P Protocol. The TLV Type field is set to 11.

6.12. Fragment



7. IPv6 Flow Specification

[[draft-raszuk-idr-flow-spec-v6](#)] defines the 12 flow specification component types for IPv6.

Type	Description
1	Destination IPv6 Prefix
2	Source IPv6 Prefix
3	Next Header
4	Port
5	Destination port
6	Source port
7	ICMP type
8	ICMP code
9	TCP flags
10	Packet length
11	Traffic Class
12	Reserved
13	Flow Label

Component Types

Each of the flow specification component types is defined in a separate TLV.

The flow specification component type is stored in the lower 8 bits of the 16-bit type field.

The format of IPv6 flow specification component type TLV that are different from the IPv4 outlined below.

7.1. Destination IPv6 Prefix TLV

Defines the destination prefix to match.

Internet-Draft OSPF Flow Specification April 2013

The IPv6 Prefix field has variable length. The length of the field is calculated from the TLV Length field.

The encoding is similar to the component type Destination IPv6 Prefix TLV. The Type field is set to 2.

The same encoding as for component type IP Protocol. The TLV Type field is set to 11.

The same encoding as for component type IP Protocol. The TLV Type field is set to 13.


In the [BGP-FLOWSPEC], the IPv4 filtering action types are standardized as BGP extended community attributes. In the OSPF, the filtering actions are carried in the TLVs of the flow-specification-LSA. The TLV type values are the same as the BGP extended community types for the filtering actions:

The traffic-action is encoded in the 2 least significant bits of the octet.

8.3. Redirect

The Redirect as defined in [[BGP-FLOWSPEC](#)] is not applicable to the OSPF.

8.4. Traffic-marking

[illegible]

The DSCP value encoded in the 6 least significant bits of the octet.

9. IPv6 Traffic Filtering Actions

[[draft-raszuk-idr-flow-spec-v6](#)] defines 3 flow specification action types. Those types are standardized as BGP extended community attributes. In OSPF, the IPv6 filtering actions are defined in TLVs. The type values are the same as the BGP extended community types for the filtering actions.

Type	Description
8006	Traffic-rate
8007	Traffic-action
8009	Traffic-marking

The flow specification filtering action type uses the full 16 bits with the highest bit set to 1.

- for inter-area and external paths, the Advertising router field of the path structure is compared.

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

b) The routing table does not have more specific unicast routes, when compared with the flow destination prefix, that have been received from a different router than the best-match unicast route, which has been determined in step a).

Under certain scenarios the validation step can be bypassed. This is outside of the scope of this document.

12. Security Considerations

As long as traffic filtering rules are restricted to match the corresponding unicast routing paths for the relevant prefixes, the security characteristics of this proposal are equivalent to the existing security properties of OSPF routing.

Where it is not the case, this would open the door to further denial-of-service attacks.

13. IANA Considerations

The following IANA assignment was made from an existing registry:

The OSPFv2 opaque LSA type TBD has been reserved for the OSPFv2 flow-specification-LSA.

Shrivastava, et al.	Expires October 19, 2013	[Page 13]
---------------------	--------------------------	-----------

Internet-Draft	OSPF Flow Specification	April 2013
----------------	-------------------------	------------

The OSPFv3 LSA function code TDB has been reserved for the OSPFv3 flow-specification-LSA.

For the purpose of this work, IANA has created a new registry entitled: "OSPF IPv4 Flow Spec Component Types". The following component types have been registered:

Type 1 - Destination Prefix

Type 2 - Source Prefix

Type 3 - IP Protocol

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

For the purpose of this work, IANA has created a new registry entitled: "OSPF IPv6 Flow Spec Component Types". The following component types have been registered:

Type 1 - Destination IPv6 Prefix

Type 2 - Source IPv6 Prefix

Type 3 - 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 - Traffic Class

Type 12 - Reserved

Type 13 - Flow Label

For the purpose of this work, IANA has created a new registry entitled: "OSPF IPv4 Flow Specification Action Types". The following component types have been registered:

Type 0x8006 - Flow spec traffic-rate

Type 0x8007 - Flow spec traffic-action

Type 0x8009 - Flow spec traffic-remarking

For the purpose of this work, IANA has created a new registry entitled: "OSPF IPv6 Flow Specification Action Types". The following component types have been registered:

Type 0x8006 - Flow spec traffic-rate

Type 0x8007 - Flow spec traffic-action

Type 0x8009 - Flow spec traffic-remarking

14. Acknowledgements

15. Normative References

[BGP-FLOWSPEC]

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

[MPLS-TE] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.

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

[OSPFV2] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.

[OSPFV3-AF]

Lindem, A., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", [RFC 5838](#), April 2010.

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

[[draft-raszuk-idr-flow-spec-v6](#)]

Raszuk, R., "Dissemination of Flow Specification Rules for IPv6", March 2012.

Authors' Addresses

Rashmi Shrivastava
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
US

Email: rashi@cisco.com

Mike Dubrovsky
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
US

Email: mdubrovs@cisco.com

Keyur Patel
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
US

Email: keyupate@cisco.com

Internet-Draft

OSPF Flow Specification

April 2013

Burjiz Pithawala
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
US

Email: bpithaw@cisco.com

