

ISIS Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 17, 2016

J. You
Q. Liang
Huawei Technologies
K. Patel
Cisco Systems
P. Fan

Z. Li
China Mobile
February 14, 2016

IS-IS Extensions for Flow Specification
draft-you-isis-flowspec-extensions-04

Abstract

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

This document discusses the use cases for distributing flow specification (FlowSpec) routes using IS-IS. Furthermore, this document defines a new IS-IS FlowSpec Reachability TLV encoding format that can be used to distribute FlowSpec rules, 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-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 17, 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 IS-IS based FlowSpec Distribution	3
3.1.	Anti-DDOS	3
4.	IS-IS Extensions for FlowSpec Rules	4
4.1.	FlowSpec Filters sub-TLV	5
4.1.1.	Order of Traffic Filtering Rules	7
4.1.2.	Validation Procedure	7
4.2.	FlowSpec Action sub-TLV	8
4.2.1.	Traffic-rate	9
4.2.2.	Traffic-action	9
4.2.3.	Traffic-marking	9
4.2.4.	Redirect-to-IP	10
5.	Redistribution of FlowSpec Rules	10
6.	IANA Considerations	11
6.1.	FlowSpec Reachability TLV	11
6.2.	FlowSpec Filters sub-TLVs	11
6.3.	FlowSpec Filter Component Types	11
6.4.	FlowSpec Action sub-TLVs	12
7.	Security Considerations	13
8.	Acknowledgement	13
9.	References	13
9.1.	Normative References	13
9.2.	Informative References	13

Authors' Addresses	14
------------------------------	--------------------

[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.

For the networks deploying only IS-IS or IS-IS variant, it is expected to extend IS-IS to distribute FlowSpec rules. This document discusses the use cases for distributing FlowSpec rules using IS-IS. Furthermore, this document also defines a new IS-IS FlowSpec Reachability TLV encoding format that can be used to distribute FlowSpec entries to the specific routers in the campus 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 IS-IS is different from BGP. The IS-IS 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 [[ISO-10589](#)] and [[RFC5575](#)].

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

[3. Use Cases for IS-IS based FlowSpec Distribution](#)

[3.1. Anti-DDOS](#)

For the networks using IS-IS or IS-IS variant, for example, the campus network or DC network, it is expected to extend IS-IS to distribute FlowSpec rules as shown in Figure 1. In this network, the traffic analyzer could be deployed to inject the FlowSpec rules into Router A. Router A creates FlowSpec entries according to the

FlowSpec rules, then the FlowSpec entries would be distributed to the other routers in this domain using IS-IS. Consequently, the attack traffic could be blocked or the suspicious traffic could be limited to a low rate as early as possible.

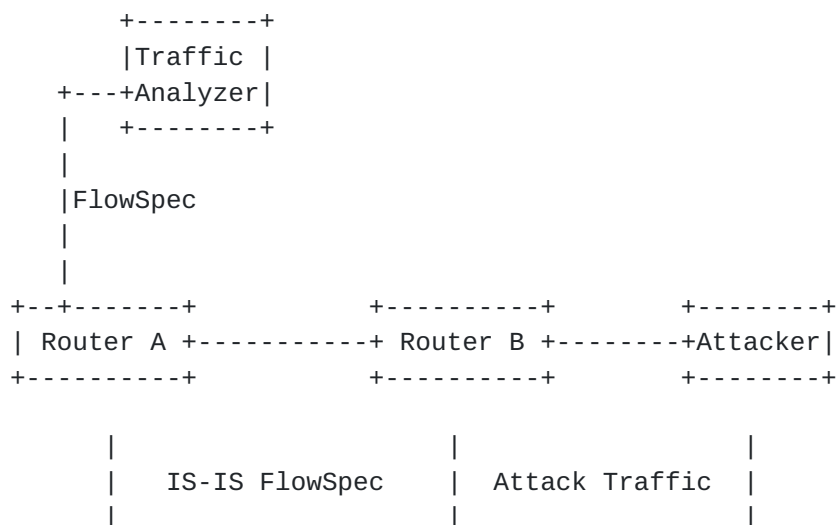


Figure 1: Anti-DDOS in IS-IS Network

4. IS-IS Extensions for FlowSpec Rules

This document defines a new IS-IS TLV, i.e. the FlowSpec reachability TLV (TLV type: TBD1), to describe the FlowSpec rules. An LSP (Link State Protocol) Data Unit [[ISO-10589](#)] can carry one or more FlowSpec reachability TLVs.

Each FlowSpec Reachability TLV carries a FlowSpec entry. The FlowSpec entry consists of a FlowSpec Filters sub-TLV and one or more corresponding FlowSpec Action sub-TLVs.

The FlowSpec Reachability TLV is defined below in Figure 2:

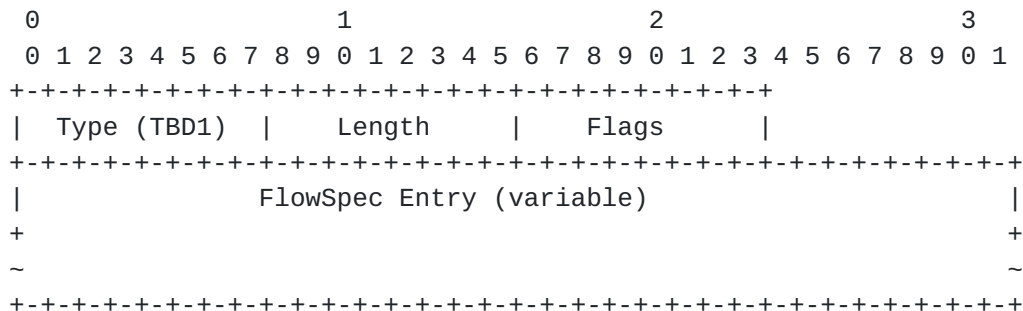


Figure 2: FlowSpec Reachability TLV

Type: 1 octet. Type code is TBD1.

Length: 1 octet. 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).

Value: variable. The value field contains a "Flags" field and a FlowSpec entry, which consists of a FlowSpec filters sub-TLV and one or more corresponding FlowSpec action sub-TLVs. The size of the FlowSpec entry cannot be greater than 253. In most scenarios, using one FlowSpec entry is sufficient. If the injected FlowSpec rule is too complex that the IS-IS router has to use more than 253 octets to encode it into a FlowSpec entry, the IS-IS router should reject it. It is strongly recommended that the FlowSpec rule provider should split or revise the complex FlowSpec rule to a suitable one for the IS-IS routers.

Flags: One octet Field identifying Flags

```

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

```

The least significant bit L is defined as a Leaking enable bit. If set, the FlowSpec Reachability TLV SHOULD be flooded across the entire routing domain. If the L flag is not set, the FlowSpec Reachability TLV MUST NOT be leaked between levels. This bit MUST NOT be altered during the TLV leaking. This Flags may be modified by the IS-IS Speaker according to a local policy.

[4.1.](#) FlowSpec Filters sub-TLV

IS-IS FlowSpec filters sub-TLV is one component of FlowSpec entry, carried in the FlowSpec reachability TLV. It is defined below in Figure 3.

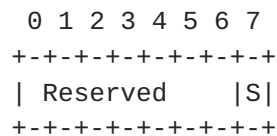


Figure 3: IS-IS FlowSpec Filters sub-TLV

Type: the TLV type (Type Code: TBD2 for IPv4 FlowSpec filters, TBD3 for IPv6 FlowSpec filters)

Length: the size of the value field in octets, it cannot be greater than 253.

Flags: One octet Field identifying Flags



The least significant bit S is defined as a strict filter check bit. If set, strict validation rules outlined in the validation section [Section 4.1.2](#) need to be enforced.

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

Table 1: IS-IS Supported FlowSpec Filter Component Types

Type	Description	RFC/ WG draft
1	Destination IPv4 Prefix	RFC5575
	Destination IPv6 Prefix	I-D.ietf-idr-flow-spec-v6
2	Source IPv4 Prefix	RFC5575
	Source IPv6 Prefix	I-D.ietf-idr-flow-spec-v6
3	IP Protocol	RFC5575
	Next Header	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

4.1.1. 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 [\[I-D.ietf-idr-flow-spec-v6\]](#).

4.1.2. 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 IS-IS FlowSpec should support similar features to mitigate the unnecessary or invalid application of

traffic filter rules. The IS-IS 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 IS-IS 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 an IS-IS speaker when the filter is considered for an installation in its FIB. An IS-IS speaker MUST flood IS-IS LSP containing a FlowSpec Reachability TLV as per the entries defined in [[ISO-10589](#)] regardless of the validation state of the prefix filters.

4.2. FlowSpec Action sub-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 IS-IS FlowSpec action TLVs, except Redirect, are same as defined in [[RFC5575](#)].


Redirect: IPv4 or IPv6 address. This target IP address MUST correspond to a tunnel in the current IS-IS router, if not, the "redirect to IP" action is invalid, and if the flowspec entry has no other action, the flowspec entry is invalid and wouldn't be installed. If the IS-IS router doesn't have a valid route for the target IP, the "redirect to IP" action is also invalid.

Table 2: BGP FlowSpec Actions

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

4.2.1. Traffic-rate


Traffic-rate TLV is encoded as:



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

4.2.2. Traffic-action

Traffic-action TLV is encoded as:



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

4.2.3. 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           TBD6           |           2           |   Reserved   | DSCP Value|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

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

4.2.4. 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           TBD7           |           6           |   Reserved   |C|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     IPv4 Address                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Redirect to IPv6 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           TBD8           |          18          |   Reserved   |C|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     IPv6 Address                                     |
|                                     |
|                                     |
|                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

IPv4/6 Address: the redirection target IP 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](#)].

5. Redistribution of FlowSpec Rules

An implementation MAY provide an option for an IS-IS speaker to announce a redistributed FlowSpec route within an IS-IS domain regardless of being installed in its local FIB. An implementation MAY impose an upper bound on number of FlowSpec entries that an IS-IS router MAY advertise.

6. IANA Considerations

This document defines the following new IS-IS TLV types, which need to be reflected in the IS-IS TLV codepoint registry.

6.1. FlowSpec Reachability TLV

Type	Description	IIH	LSP	SNP
TBD1	The FlowSpec Reachability TLV	n	y	n

6.2. FlowSpec Filters sub-TLVs

Type	Description
TBD2	IPv4 FlowSpec filters sub-TLV
TBD3	IPv6 FlowSpec filters sub-TLV

6.3. FlowSpec Filter Component Types

Type	Description	RFC/ WG draft
1	Destination IPv4 Prefix	RFC5575
	Destination IPv6 Prefix	I-D.ietf-idr-flow-spec-v6
2	Source IPv4 Prefix	RFC5575
	Source IPv6 Prefix	I-D.ietf-idr-flow-spec-v6
3	IP Protocol	RFC5575
	Next Header	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

6.4. FlowSpec Action sub-TLVs

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

Type TBD4 - traffic-rate

Type TBD5 - traffic-action

Type TBD6 - traffic-marking

Type TBD7 - redirect to IPv4

Type TBD8 - redirect to IPv6

7. Security Considerations

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

8. Acknowledgement

The authors would like to thank Jeff Haas for his useful comments.

9. References

9.1. Normative References

[ISO-10589]

ISO, "Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", International Standard 10589: 2002, Second Edition, 2002.

[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>>.

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

[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-02](#) (work in progress), January 2014.

`[I-D.ietf-idr-flow-spec-v6]`

Raszuk, R., Pithawala, B., McPherson, D., and A. Andy,
"Dissemination of Flow Specification Rules for IPv6",
[draft-ietf-idr-flow-spec-v6-06](#) (work in progress),
November 2014.

`[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.ietf-idr-flowspec-redirect-rt-bis]`

Haas, J., "Clarification of the Flowspec Redirect Extended
Community", [draft-ietf-idr-flowspec-redirect-rt-bis-05](#)
(work in progress), July 2015.

Authors' Addresses

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

Email: youjianjie@huawei.com

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

Email: liangqiandeng@huawei.com

Keyur Patel
Cisco Systems
170 W. Tasman Drive
San Jose CA 95124 95134
US

Email: keyupate@cisco.com

Peng Fan
Beijing
China

Email: peng.fan@139.com

Zhenqiang Li
China Mobile
Beijing
China

Email: li_zhenqiang@hotmail.com