Isis Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 31, 2015

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

June 29, 2015

# IS-IS Extensions for Flow Specification
# draft-you-isis-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., IS-IS), 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 IS-IS. Furthermore, this document defines a new IS-IS FlowSpec Reachability TLV 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-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 December 31, 2015.



Table of Contents

## 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., IS-IS), it is expected to extend the IGP (IS-IS in this document) to distribute FlowSpec routes. This document discusses the use cases for distributing FlowSpec routes using IS-IS. Furthermore, this document also defines a new IS-IS FlowSpec Reachability TLV 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 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, including filters and actions. Each FlowSpec consists of a set of filters and a set of actions.

## 3. Use Cases for IS-IS based FlowSpec Distribution

For the networks deploying only an IGP (e.g., IS-IS), it is expected to extend the IGP (IS-IS 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 IS-IS based FlowSpec route distribution.

### 3.1. IS-IS Campus Network

For networks not deploying BGP, for example, the campus network using IS-IS, it is expected to extend IS-IS to distribute FlowSpec routes as shown in Figure 1. 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 IS-IS.

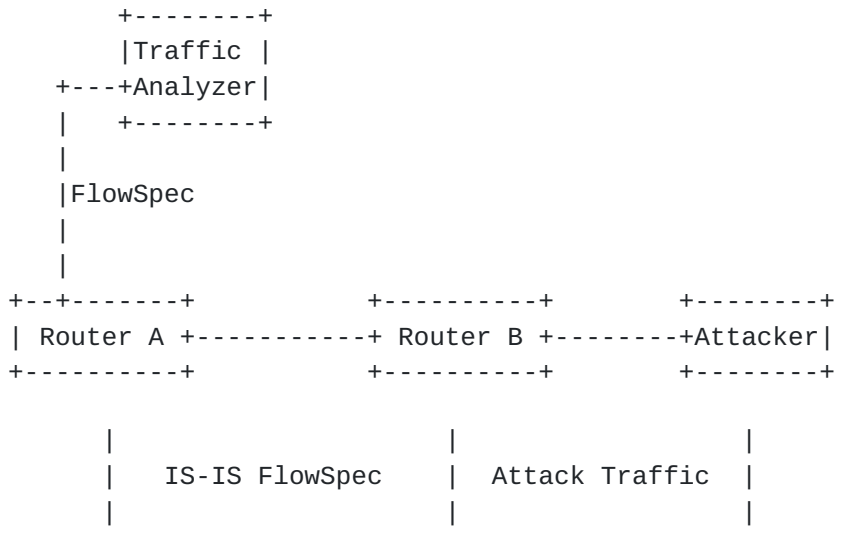


Figure 1: IS-IS Campus Network

### 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 2. 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 IS-IS, 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.

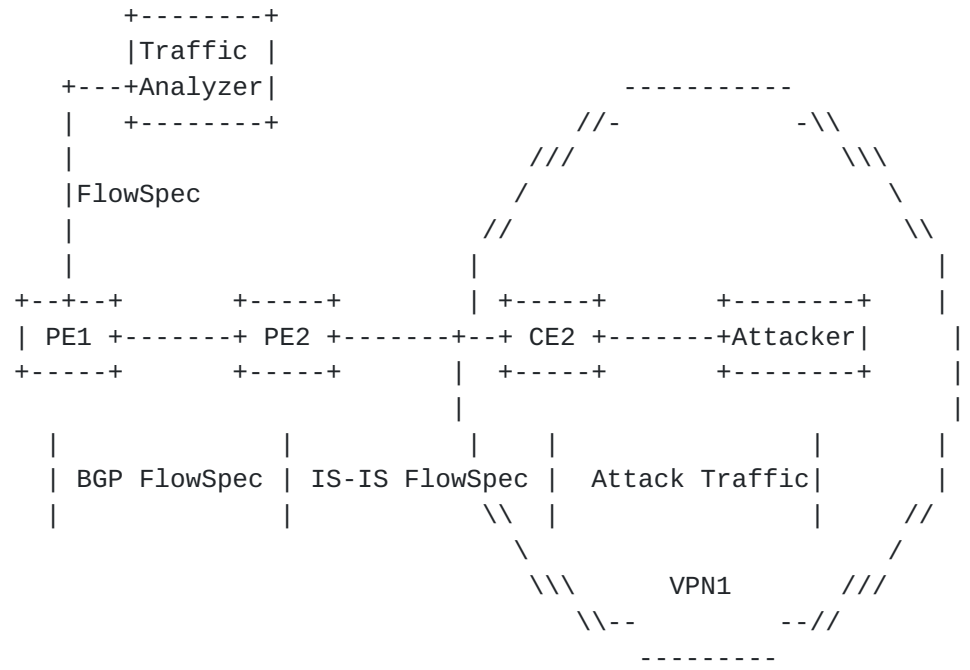


Figure 2: Traffic Analyzer deployed in Provider Network

### 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 3. 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 IS-IS 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.

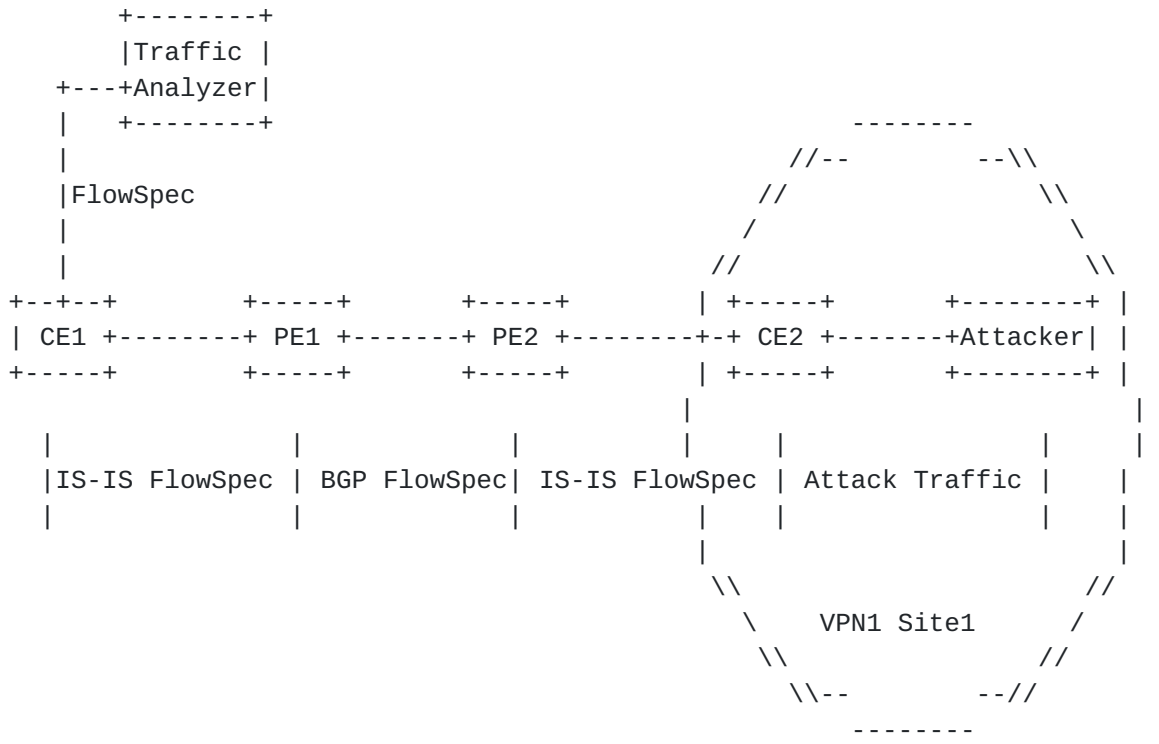


Figure 3: Traffic Analyzer deployed in Customer Network

### 3.2.3. Policy Configuration

The CE or PE could deploy local filtering policies to filter IS-IS FlowSpec rules, for example, deploying a filtering policy to filter the incoming IS-IS 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 IS-IS Instance FlowSpec RIB. Otherwise, the PE couldn't transform a BGP

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

## 4. IS-IS Extensions for FlowSpec Routes

This document defines a new IS-IS TLV, i.e. the FlowSpec reachability TLV (TLV type: TBD1), which would be carried in an LSP (Link State Protocol) Data Unit [ISO-10589], to describe the FlowSpec routes.

The FlowSpec Reachability TLV carries one or more FlowSpec entries. Each FlowSpec entry consists of FlowSpec filters (FlowSpec filters sub-TLVs) and corresponding FlowSpec actions (FlowSpec Action sub-TLVs).

The FlowSpec Reachability TLV is defined below in Figure 4:

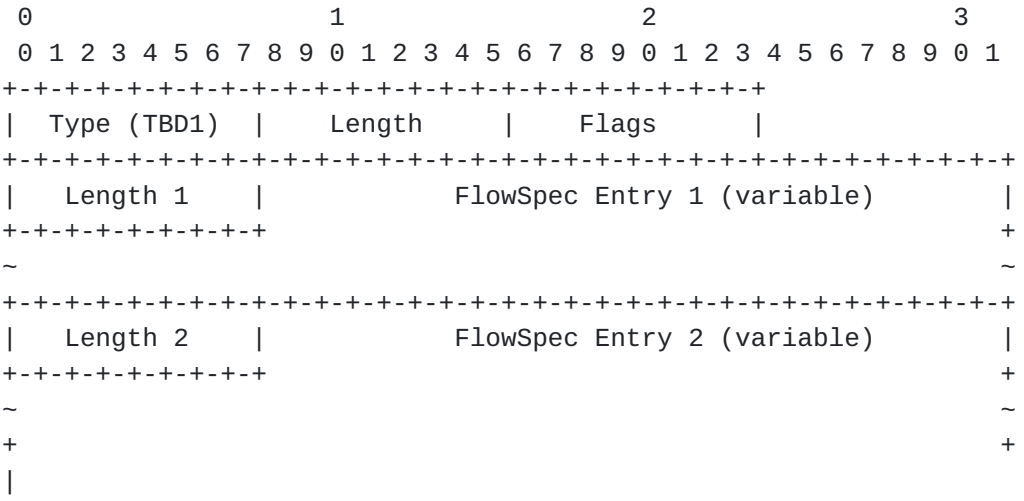


Figure 4: FlowSpec Reachability TLV

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 one or more 2-tuples consisting of the Length and the FlowSpec entry. Each 2-tuple starts with 1 octet of Length, and followed by a variable length FlowSpec entry, which consists of FlowSpec filters sub-TLVs and corresponding FlowSpec action sub-TLVs. The length specifies the number of bytes of the FlowSpec entry.

Flags: One octet Field identifying Flags

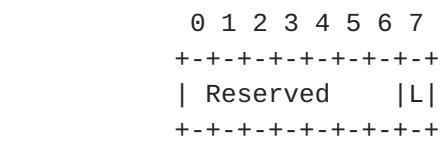


The least significant bit L is defined as a Leaking enable bit. If set, the FlowSpec Reachability TLV SHOULD be flooded across the entire routing domain. If the L flag is not set, the FlowSpec Reachability TLV MUST NOT be leaked between levels. This bit MUST NOT be altered during the TLV leaking. This 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 5.

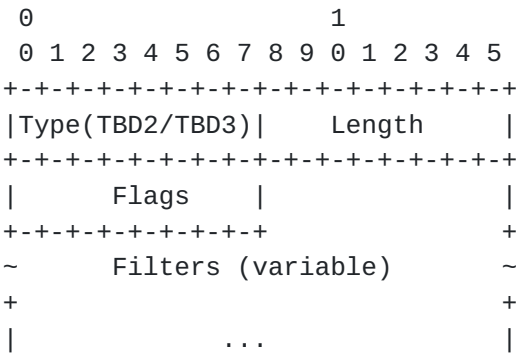


Figure 5: IS-IS FlowSpec Filters sub-TLV

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

Flags: One octet Field identifying Flags

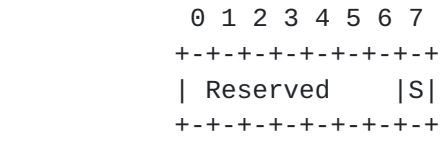


The least significant bit S is defined as a strict filter check bit. If set, strict validation rules outlined in the validation section Section 4.1.2 need to be enforced.

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

Table 1: IS-IS Supported FlowSpec Filters

| 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 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 rules 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 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 |

### 4.2.1. Traffic-rate

Traffic-rate TLV is encoded as:

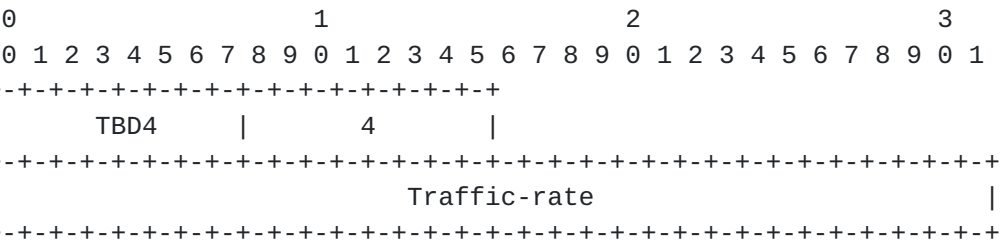


Traffic-rate: the same as defined in [RFC5575].

### 4.2.2. Traffic-action

Traffic-action TLV is encoded as:

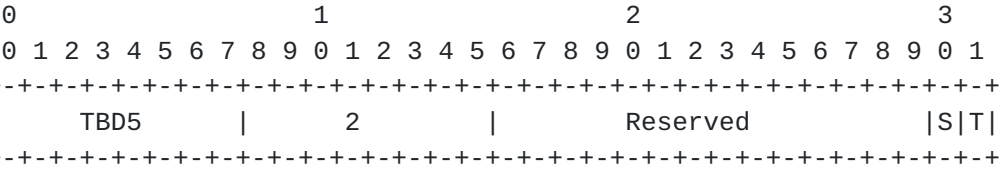


S flag and T flag: the same as defined in [RFC5575].

### 4.2.3. Traffic-marking

Traffic-marking TLV is encoded as:

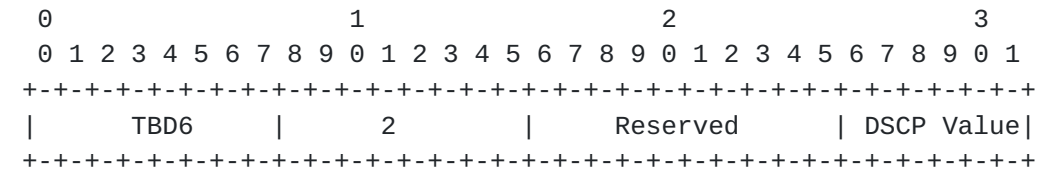


DSCP value: the same as defined in [RFC5575].

### 4.2.4. Redirect-to-IP

Redirect-to-IPv4 is encoded as:

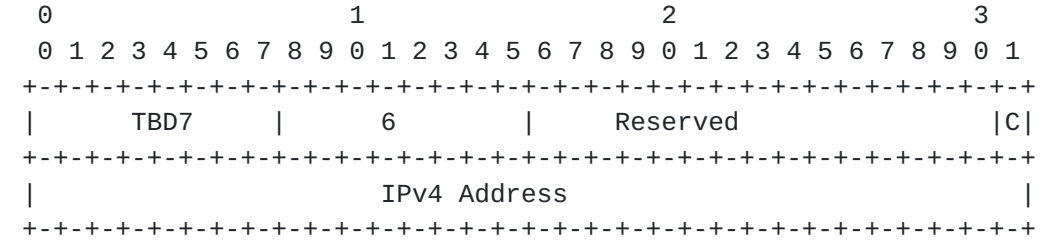


Redirect to IPv6 TLV is encoded as:

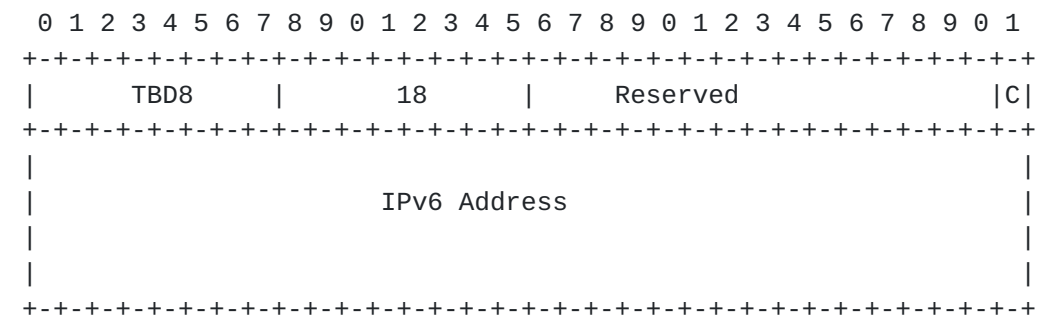


IPv4/6 Address: the redirection target address.

'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].

## 5. Redistribution of FlowSpec Routes

In certain scenarios, FlowSpec routes MAY get redistributed from one protocol domain to another; specifically from BGP to IS-IS and vice-versa. When redistributed from BGP, the IS-IS speaker SHOULD

generate a FlowSpec Reachability TLV for the redistributed routes and announce it within an IS-IS domain. 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 routes 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-TLV

| Type | Description | encoding |
|---|---|---|
| TBD2<br>TBD3 | The FlowSpec filters sub-TLV | flow-spec NLRI value [RFC5575]<br>I-D.ietf-idr-flow-spec-v6 |

### 6.3. FlowSpec Action sub-TLV

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

TBD.

## 9. References

### 9.1. Normative References

[ISO-10589] ISO, "Intermediate System to Intermediate System intra-domain routeing 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, March 1997.

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

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

### 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-04 (work in progress), April 2015.

## Authors' Addresses

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

Email: youjianjie@huawei.com

Qiandeng Liang
Huawei
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
USA

Email: keyupate@cisco.com

Peng Fan
China Mobile

Email: fanpeng@chinamobile.com