

IDR  
Internet-Draft  
Updates: [5575](#) (if approved)  
Intended status: Standards Track  
Expires: September 4, 2018

Z. Li  
China Mobile  
J. Dong  
S. Zhuang  
Huawei Technologies  
March 3, 2018

**Populate to FIB Action for FlowSpec  
draft-li-idr-flowspec-populate-to-fib-02**

Abstract

A bit, F bit, is defined in traffic action extended community, which is used by FlowSpec to indicate the associated specifications be populated in FIB (Forwarding Information Base) after appropriate process.

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 <https://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 September 4, 2018.

Copyright Notice

Copyright (c) 2018 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 (<https://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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Requirements Language . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Populate to FIB Action . . . . .	<a href="#">3</a>
<a href="#">4.</a>	Implementation Considerations . . . . .	<a href="#">3</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">4</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">4</a>
<a href="#">7.</a>	Normative References . . . . .	<a href="#">4</a>
	Authors' Addresses . . . . .	<a href="#">5</a>

## [1.](#) Introduction

BGP FlowSpec [[RFC5575](#)] provides a flexible mechanism to distribute traffic flow specifications, where the matching rules are encoded in the Border Gateway Protocol Network Layer Reachability Information (BGP NLRI) with defined new format and the corresponding actions are encoded in BGP Extended communities.

In routers, traffic flow specifications distributed by BGP FlowSpec [[RFC5575](#)] are stored in distinct set of RIBs (Routing Information Base) according to their (AFI, SAFI) pairs. These RIBs are then populated to the dedicated hardware (most of them are TCAM based) usually shared with ACLs (Access Control Lists). The dedicated hardware is much more expensive and space limited when compared with the hardware used to store the FIB (Forwarding Information Base), which is usually sufficient to fit several millions of FIB entries. Although in some implementations, the hardware used to populate traffic flow specifications and FIB entries is the same, the size for each parts is fixed at design stage. As the number of ACL rules and FlowSpec specifications increases, especially when FlowSpec is used for dynamic traffic flow steering, which is one of the three BGP FlowSpec applications listed in [[RFC5575](#)] and [[I-D.ietf-idr-rfc5575bis](#)], hardware space requirement of FlowSpec specifications in the field network may exceed the size of the dedicated hardware. To save the limited and expensive space of the dedicated hardware, it is better to populate some FlowSpec specifications to FIB if possible. The destination prefix based FlowSpec specifications, for example, are suitable to be populated to FIB.

However, there is no method in the current version of BGP FlowSpec [[RFC5575](#)] and RFC5575bis [[I-D.ietf-idr-rfc5575bis](#)] to indicate the associated specifications are suitable to be populated to FIB. This



document defines a new bit, F bit (populate to FIB), in 0x8007 traffic action extended community to satisfy the requirement.

## 2. 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 [\[RFC2119\]](#).

## 3. Populate to FIB Action

F bit, populate to FIB bit, is defined in 0x8007 traffic action extended community [\[RFC5575\]](#) to indicate the associated BGP FlowSpec specifications are suitable to be populated to FIB. Thus the space of the dedicated hardware that is used to store the BGP FlowSpec specifications can be saved for other kinds of BGP FlowSpec specifications and ACL rules.

The encoding format of the traffic action extended community with F bit is shown below. The F bit is solicited to be assigned by IANA.

```

    40  41  42  43  44  45  46  47
    +---+---+---+---+---+---+---+
    |           reserved       | F | S | T |
    +---+---+---+---+---+---+---+

```

Traffic-action extended community consists of 2 bytes for type and subtype, the value of which MUST be 0x8007, and 6 bytes for value, of which only the 3 least significant bits of the 6th byte (from left to right) are currently defined. S and T are defined in BGP FlowSpec [\[RFC5575\]](#). F is defined as:

- o F: Populate to FIB Action (bit 45, to be assigned by IANA): When this bit is set, the associated BGP FlowSpec specifications SHOULD be populated to FIB. If not set, the associated BGP FlowSpec specifications MUST NOT be populated to FIB. If this bit is set and the associated BGP FlowSpec specifications can not be populated to FIB, the associated BGP FlowSpec specifications MUST be ignored.

## 4. Implementation Considerations

FlowSpec rules are ordering sensitive. After ordering processing as per [section 5.1 of \[RFC5575\]](#), they are searched sequentially until a matching rule is found. FIB entries, on the contrary, have no ordering implication. Longest prefix matching is the rule to choose the matching FIB entry. Only the destination prefix based, F bit tagged FlowSpec rules that pass the validation (as per [section 6](#) of



[RFC5575]) and ordering (as per [section 5.1 of \[RFC5575\]](#)) processing are suitable to be populated into FIB. When populating a FlowSpec rule into FIB, the following facts have to be taken into account.

- o FlowSpec rules have higher priority than corresponding IGP and BGP routing entries.
- o When populating the FIB, the FlowSpec rules with F bit tagged are preferred than the corresponding IGP and BGP routing entries.
- o When a FlowSpec rule is being populated into FIB, the FIB entries, including those come from IGP or BGP updates, covered by this FlowSpec rule MUST be removed or replaced by this FlowSpec rule.
- o The populated FlowSpec rules in the FIB MUST not be overridden by IGP or BGP updates.

## 5. Security Considerations

This document defines a new bit in the traffic action extended community to indicate the associated BGP FlowSpec specifications SHOULD be populated to FIB directly. This bit does not introduce any new security issues. The same security considerations as for the BGP FlowSpec [\[RFC5575\]](#) applies.

## 6. IANA Considerations

One bit, F bit, is solicited to be assigned from Traffic Action Fields registry. This bit is used by BGP FlowSpec to indicate the associated BGP FlowSpec specifications SHOULD be populated to FIB directly.

## 7. Normative References

[I-D.ietf-idr-rfc5575bis]

Hares, S., Loibl, C., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", [draft-ietf-idr-rfc5575bis-06](#) (work in progress), October 2017.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.



[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](https://www.rfc-editor.org/info/rfc5575), DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.

#### Authors' Addresses

Zhenqiang Li  
China Mobile  
No.32 Xuanwumenxi Ave., Xicheng District  
Beijing 100032  
P.R. China

Email: [li\\_zhenqiang@hotmail.com](mailto:li_zhenqiang@hotmail.com)

Jie Dong  
Huawei Technologies  
Huawei Campus, No. 156 Beiqing Rd.  
Beijing 100095  
China

Email: [jie.dong@huawei.com](mailto:jie.dong@huawei.com)

Shunwan Zhuang  
Huawei Technologies  
Huawei Campus, No. 156 Beiqing Rd.  
Beijing 100095  
China

Email: [zhuangshunwan@huawei.com](mailto:zhuangshunwan@huawei.com)



