

IDR
Internet-Draft
Intended status: Informational
Expires: January 3, 2016

G. Van de Velde
W. Henderickx
Alcatel-Lucent
N. Fevrier
A. Karch
Cisco Systems
A. Grewal
Juniper Networks
July 2, 2015

Dissemination of Flow Specification Rules for IPv6 Implementation Report
[draft-vandvelde-idr-ipv6-flowspec-imp-02](#)

Abstract

This document is an implementation report for the BGP Flow Specification Rules for IPv6 as defined in [[I-D.ietf-idr-flow-spec-v6](#)]. The respondents are experts with the implementations they reported on, and their responses are considered authoritative for the implementations for which their responses represent.

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 January 3, 2016.

Copyright Notice

Copyright (c) 2015 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

Internet-Draft

July 2015

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	2
2.	Requirements Language	2
3.	Implementation Forms	3
4.	NLRI and Extended Community subtypes	3
5.	Interoperable Implementations	6
5.1.	Cisco Systems - Alcatel-Lucent	6
5.2.	Cisco Systems - Juniper Networks	8
5.3.	Juniper Networks - Cisco Systems	10
6.	IANA Considerations	12
7.	Security Considerations	12
8.	Privacy Considerations	12
9.	Acknowledgements	13
10.	Change Log	13
11.	References	13
11.1.	Normative References	13
11.2.	Informative References	13
	Authors' Addresses	13

[1.](#) Introduction

In order to share Flow Specification Rules for IPv6 using the BGP routing protocol a new BGP Network Layer Reachability Information (NLRI) encoding format is required.

This document provides an implementation report for the BGP Dissemination of Flow Specification Rules for IPv6 NLRI Format as defined in [[I-D.ietf-idr-flow-spec-v6](#)].

The editors did not verify the accuracy of the information provided by respondents or by any alternative means. The respondents are experts with the implementations they reported on, and their responses are considered authoritative for the implementations for which their responses represent.

[2.](#) Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [[RFC2119](#)] only when they appear in all

upper case. They may also appear in lower or mixed case as English words, without any normative meaning.

[3.](#) Implementation Forms

Contact and implementation information for person filling out this form:

Cisco

Name: Nicolas Fevrier
Email: nifevrie@cisco.com
Vendor: Cisco Systems, Inc.
Release: IOS-XR 5.3.1
Protocol Role: Sender, Receiver

Alcatel-Lucent

Name: Wim Henderickx
Email: wim.henderickx@alcatel-lucent.com
Vendor: Alcatel-Lucent, Inc.
Release: R12R4
Protocol Role: Sender, Receiver

Juniper

Name: Ashutosh Grewal
Email: agrewal@juniper.net
Vendor: Juniper Networks, Inc.
Release: JunOS 15.2
Protocol Role: Sender, Receiver

[4.](#) NLRI and Extended Community subtypes

Does the implementation support the Network Layer Reachability (NLRI) subtypes as described in [Section 3](#) and 4 of [[I-D.ietf-idr-flow-spec-v6](#)].

- o N1: Type 1 - Destination IPv6 Prefix

- o N2: Type 2 - Source IPv6 Prefix
- o N3: Type 3 - Next Header
- o N4: Type 4 - Port
- o N5: Type 5 - Destination port
- o N6: Type 6 - Source port
- o N7: Type 7 - ICMP type

- o N8: Type 8 - ICMP code
- o N9: Type 9 - TCP flags
- o N10: Type 10 - Packet length
- o N11: Type 11 - DSCP (Diffserv Code Point)
- o N12: Type 12 - Fragment
- o N13: Type 13 - Flow Label
- o E1: Extended Community - traffic-rate
- o E2: Extended Community - traffic-action
- o E3: Extended Community - redirect
- o E4: Extended Community - traffic-marking

Internet-Draft

July 2015

		Cisco	ALU	Juniper
Rcv.N1	YES	YES	YES	YES
Snd.N1	YES	YES	YES	YES
Rcv.N2	YES	YES	YES	YES
Snd.N2	YES	YES	YES	YES
Rcv.N3	YES	YES	YES	YES
Snd.N3	YES	YES	YES	YES
Rcv.N4	YES	YES	YES	YES
Snd.N4	YES	YES	YES	YES
Rcv.N5	YES	YES	YES	YES
Snd.N5	YES	YES	YES	YES
Rcv.N6	YES	YES	YES	YES
Snd.N6	YES	YES	YES	YES
Rcv.N7	YES	YES	YES	YES
Snd.N7	YES	YES	YES	YES
Rcv.N8	YES	YES	YES	YES
Snd.N8	YES	YES	YES	YES
Rcv.N9	YES	YES	YES	YES
Snd.N9	YES	YES	YES	YES
Rcv.N10	YES	YES	YES	YES

Snd.N10	YES	YES	YES	
Rcv.N11	YES	YES	YES	
Snd.N11	YES	YES	YES	
Rcv.N12	YES	YES	YES	
Snd.N12	YES	YES	YES	
Rcv.N13	YES	YES	YES	
Snd.N13	YES	YES	YES	
Rcv.E1	YES	YES	YES	
Snd.E1	YES	YES	YES	
Rcv.E2	YES	YES	YES	
Snd.E2	YES	YES	YES	
Rcv.E3	YES	YES	YES	
Snd.E3	YES	YES	YES	
Rcv.E4	YES	YES	YES	
Snd.E4	YES	YES	YES	
+-----+-----+-----+-----+				

Yes

- o Rcv: BGP speaker can receive the information into the BGP process
- o Snd: BGP speaker can relay the information from the BGP process

No

- o Rcv: BGP speaker cannot receive the information into the BGP process
- o Snd: BGP speaker can not relay the information from the BGP process

[5.](#) Interoperable Implementations

Summary of executed Interop tests between different implementations

[5.1.](#) Cisco Systems - Alcatel-Lucent

This Interop test was between a Cisco router and a Alcatel-Lucent router. Between the two BGP devices an iBGP session is established.

The following IPv6 Flow Specification NLRI is constructed using the Cisco router as IPv6 Flow Specification controller:

```
!  
class-map type traffic match-all InteropMatchList  
  match destination-address ipv6 2001:2::3/128  
  match source-address ipv6 2002:2::3/128  
  match destination-port 1-5 7-11 13-18 20-25 27-31  
  match source-port 33-37 39-43 45-50 53-58 60-65  
  match ipv6 icmp-type 35  
  match ipv6 icmp-code 55
```



```

BGP Router ID:195.207.5.200    AS:65117    Local AS:65117
=====
Legend -
Status codes   : u - used, s - suppressed, h - history, d - decayed, * - valid
                  l - leaked, x - stale, > - best, b - backup
Origin codes   : i - IGP, e - EGP, ? - incomplete
=====
BGP FLOW IPV6 Routes
=====
Flag  Network                Nexthop                LocalPref    MED
      As-Path
-----
u*>i  --                    2001:::1          100          None
      No As-Path

Community Action: ext:800:0
Community Action: rate-limit: 65117:1103626240
Community Action: mark-dscp: 45
NLRI Subcomponents:
Dest Pref  : 2001:2::3/128 offset 0
Src Pref   : 2002:2::3/128 offset 0
Ip Proto   : [ >= 6 ] and [ <= 71 ] or [ >= 73 ] and [ <= 80 ] or [ >=
Dest Port  : [ >= 1 ] and [ <= 5 ] or [ >= 7 ] and [ <= 11 ] or [ >= 13
Src Port   : [ >= 33 ] and [ <= 37 ] or [ >= 39 ] and [ <= 43 ] or [ >=
ICMP Type  : [ == 35 ]
ICMP Code  : [ == 55 ]
TCP Flags  : [ 240 ]
TCP Flags  : [ 240 ]
DSCP       : [ >= 1 ] and [ <= 10 ] or [ >= 11 ] and [ <= 20 ] or [ >=
Frag       : [ == 4 ]
=====

```

```

Routes : 1
=====

```

[5.2.](#) Cisco Systems - Juniper Networks

This Interop test was between a Cisco router and a Juniper router. Between the two BGP devices an iBGP session is established.

The following IPv6 Flow Specification NLRI is constructed using the Cisco router as IPv6 Flow Specification controller:

This results with the following Flow Specification Extended communities and IPv6 Flow Specification NLRI:

The above IPv6 Flow Specification rule is correctly received by the

Juniper BGP speaker and is reflected as follows on the device:

Internet-Draft

July 2015

```
root@sdn-st-mx480-b> show route table inet6flow.0 extensive
inet6flow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
2001:2::3/128,2002:2::3/128,proto>=6<=71,>=73<=80,>=85<=90,>=95<=105,>=110<=115,
dstport>=1<=5,>=7<=11,>=13<=18,>=20<=25,>=27<=31,
srcport>=33<=37,>=39<=43,>=45<=50,>=53<=58,>=60<=65,icmp6-type=35,icmp6-co
tcp-flag:f0,len>=120<=130,>=135<=140,>=145<=160,>=165<=200,>=205<=225,
dscp>=1<=10,>=11<=20,>=22<=30,>=32<=40,>=52<=60,frag=04/term:N/A
(1 entry, 0 announced)
  *BGP      Preference: 170/-101
            Next hop type: Fictitious, Next hop index: 0
            Address: 0x95542a4
            Next-hop reference count: 2
            State: <Active Int Ext>
            Local AS: 65117 Peer AS: 65117
            Age: 1:58
            Validation State: unverified
            Task: BGP_65117.10.0.0.2
            AS path: I
            Communities: unknown iana 800 traffic-rate:65117:25 traffic-mar
            Accepted
            Localpref: 100
            Router ID: 10.0.0.2
Jun 12 15:37:57.990156 BGP RECV 10.0.0.2+179 -> 10.0.0.1+58360
Jun 12 15:37:57.990202 BGP RECV message type 2 (Update) length 245
Jun 12 15:37:57.990228 BGP RECV Update PDU length 245
Jun 12 15:37:57.990260 BGP RECV flags 0x90 code MP_reach(14): AFI/SAFI 2/133
Jun 12 15:37:57.990287 -akg- nhlen from packet : 16
Jun 12 15:37:57.990316 BGP RECV          nhop zero-len len 16
Jun 12 15:37:57.990388 BGP RECV          2001:2::3/128,2002:2::3/128,proto>=6<=
>=85<=90,>=95<=105,>=110<=115,dstport>=1<=5,>=7<=11,>=13<=18,>=20<=25,>=
srcport>=33<=37,>=39<=43,>=45<=50,>=53<=58,>=60<=65,icmp6-type=35,icmp6-co
tcp-flag:f0,len>=120<=130,>=135<=140,>=145<=160,>=165<=200,>=205<=225,
dscp>=1<=10,>=11<=20,>=22<=30,>=32<=40,>=52<=60,frag=04/1240
Jun 12 15:37:57.990418 BGP RECV flags 0x40 code Origin(1): IGP
Jun 12 15:37:57.990446 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Jun 12 15:37:57.990471 BGP RECV flags 0x40 code LocalPref(5): 100
Jun 12 15:37:57.990517 BGP RECV flags 0xc0 code Extended Communities(16): <snip
```

[5.3.](#) Juniper Networks - Cisco Systems

This Interop test was between a Juniper router and a Cisco router.
Between the two BGP devices an iBGP session is established.

The following IPv6 Flow Specification NLRI is constructed using the
Juniper router as IPv6 Flow Specification controller:

```
!
root@sdn-st-mx480-b# show routing-options
rib inet6.0 {
  flow {
    route flowroute {
      match {
        destination abcd::1/128;
        source abcd::2/128;
        protocol [ 6-71 73-80 85-90 95-105 110-115 ];
        destination-port [ 1-5 7-11 13-18 20-25 27-31 ];
        source-port [ 33-37 39-43 45-50 53-58 60-65 ];
        icmp6-type 35;
        icmp6-code 55;
        tcp-flags 240;
        packet-length [ 120-130 135-140 145-160 165-200 205-225 ];
        dscp [ 1-10 11-20 22-30 32-40 52-60 ];
        fragment first-fragment;
        flow-label [ ( 22 || 77 ) 11 33 89 ];
      }
      then {
        rate-limit 9600;
        mark 45;
      }
    }
  }
}
```

The above IPv6 Flow Specification rule is correctly received by the
Cisco BGP speaker and is reflected as follows on the device:

July 2015

6. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: The IANA has requested that this section remain in the document upon publication as an RFC. This note to the RFC Editor, however, may be removed.

7. Security Considerations

No new security issues are introduced to the BGP defined in Dissemination of Flow Specification Rules for IPv6 [[I-D.ietf-idr-flow-spec-v6](#)].

8. Privacy Considerations

No new privacy issues are introduced to the BGP defined in Dissemination of Flow Specification Rules for IPv6 [[I-D.ietf-idr-flow-spec-v6](#)].

Van de Velde, et al. Expires January 3, 2016 [Page 12]

Internet-Draft

July 2015

9. Acknowledgements

The authors would like to thank Hyojeong Kim, Bertrand Duvivier, R. Divya, and Adam Simpson.

10. Change Log

Initial Version: 8 October 2014

-01 Version: 20 May 2015

-02 Version: 2 July 2015

11. References

11.1. Normative References

[[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.

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

[11.2](#). Informative References

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway
Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

Authors' Addresses

Gunter Van de Velde
Alcatel-Lucent
Copernicuslaan 50
Antwerpen 2018
Belgium

Email: gunter.van_de_velde@alcatel-lucent.com

Wim Henderickx
Alcatel-Lucent

Email: wim.henderickx@alcatel-lucent.be

Van de Velde, et al. Expires January 3, 2016

[Page 13]

Internet-Draft

July 2015

Nicolas Fevrier
Cisco Systems
11 Rue Camille Desmoulins
Issy-les-Koulineaux 92130
France

Email: nifevrie@cisco.com

Andy Karch
Cisco Systems
170 W. Tasman Drive

San Jose, CA 95124 95134
USA

Email: akarch@cisco.com

Ashutosh Grewal
Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089 94089
USA

Email: agrewal@juniper.net