

Network Working Group
Internet Draft
Updates: [5575](#)
Intended Status: Proposed Standard
Expiration Date: December 2012

James Uttaro
AT&T
Clarence Filselfil
Pradosh Mohapatra
David Smith
Cisco
June 25, 2012

Revised Validation Procedure for BGP Flow Specifications
draft-ietf-idr-bgp-flowspec-oid-00

Abstract

This document describes a modification to the validation procedure defined in [RFC 5575](#) for the dissemination of BGP flow specifications. [RFC 5575](#) requires that the originator of the flow specification matches the originator of the best-match unicast route for the destination prefix embedded in the flow specification. This allows only BGP speakers within the data forwarding path (such as autonomous system border routers) to originate BGP flow specifications. Though it is possible to disseminate such flow specifications directly from border routers, it may be operationally cumbersome in an autonomous system with a large number of border routers having complex BGP policies. The modification proposed herein enables flow specifications to be originated from a centralized BGP route controller.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on December 1, 2012.

Copyright Notice

Copyright (c) 2012 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	Specification of Requirements	2
2	Motivation	3
3	Introduction	5
4	Revised Validation Procedure	6
5	Security Considerations	7
6	IANA Considerations	7
7	Normative References	7
8	Acknowledgements	8
9	Authors' Addresses	8

[1](#). Specification of Requirements

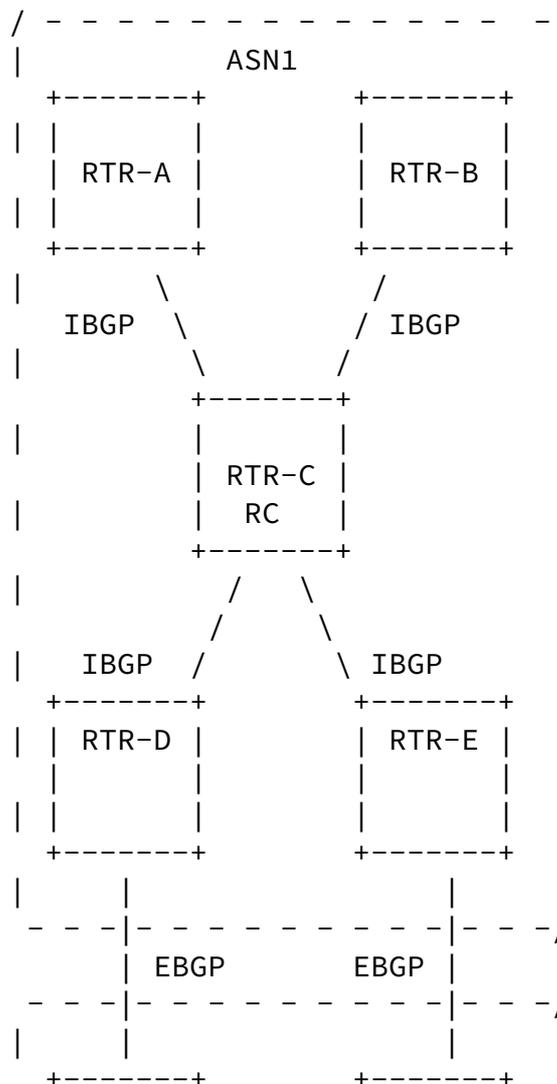
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this

document are to be interpreted as described in [[RFC2119](#)].

[2](#). Motivation

Step (a) of the validation procedure in [[RFC5575](#)], [section 6](#) is defined with the underlying assumption that the flow specification NLRI traverses the same path, in the inter-domain and intra-domain route distribution graph, as that of the longest-match unicast route for the destination prefix embedded in the flow specification.

In the case of inter-domain traffic filtering, for example, the flow specification originator at the egress border routers of ASN1 (RTR-D and RTR-E in figure 1) matches the EBGP neighbor that advertised the longest match destination prefix (RTR-F and RTR-G respectively). Similarly, at the ingress border routers of ASN1 (RTR-A and RTR-B in figure 1), the flow specification originator matches the egress IBGP border routers that had advertised the unicast route for the best-match destination prefix (RTR-D and RTR-E respectively). This is true even when ingress border routers select paths from different egress border routers as best path based upon IGP distance (as an example, RTR-A chooses RTR-D's path as best; RTR-B chooses RTR-E as the best path).



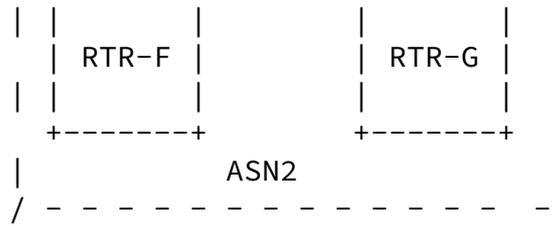


Figure 1

It is highly desirable that each ASN is able to protect itself independently from network security attacks using the BGP flow specification NLRI for intra-domain purposes only. Network operators often deploy a dedicated Security Operations Center (SOC) within their ASN to monitor and detect such security attacks. To mitigate attacks in a scalable intra-domain manner, operators require the ability to originate intra-domain flow specification NLRIs from a central BGP route controller (or router reflector per [RFC4456]) that is not within the data forwarding plane. In this way, operators can direct border routers within their ASN with specific attack

mitigation actions (drop the traffic, forward to a clean-pipe center, etc.). To originate a flow specification NLRI, a central BGP route controller (or route reflector) must set itself as the originator in the flowspec NLRI. This is necessary given the route controller is originating the flow specification not reflecting it, and to avoid the complexity of having to determine the egress border router whose path was chosen as the best in each of the ingress border routers. It thus becomes necessary to modify step (a) of the [RFC 5575](#) validation procedure such that an IBGP peer that is not within the data forwarding plane may originate flow specification NLRIs.

3. Introduction

[RFC 5575](#) defined a new BGP capability that can be used to distribute traffic flow specifications amongst BGP speakers in support of traffic filtering. The primary intention of [RFC 5575](#) is to enable downstream autonomous systems to signal traffic filtering policies to upstream autonomous systems. In this way, traffic is filtered closer to the source and the upstream autonomous system(s) avoid carrying the traffic to the downstream autonomous system only to be discarded. [RFC 5575](#) also enables more granular traffic filtering based upon

upper layer protocol information (e.g., protocol port numbers) as opposed to coarse IP destination prefix-based filtering. Flow specification NLRIs received from a BGP peer are subject to validity checks before being considered feasible and subsequently installed within the respective Adj-RIB-In. The validation procedure defined within [RFC 5575](#) requires that the originator of the flow specification NLRI matches the originator of the best-match unicast route for the destination prefix embedded in the flow specification. This allows only BGP speakers [[RFC4271](#)] within the data forwarding path (such as autonomous system border routers) to originate BGP flow specification NLRIs. Though it is possible to disseminate such flow specification NLRIs directly from border routers, it may be operationally cumbersome in an autonomous system with a large number of border routers having complex BGP policies. This document describes a modification to the [RFC 5575](#) validation procedure allowing flow specification NLRIs to be originated from a centralized BGP route controller within the local autonomous system that is neither in the data forwarding path nor serving as a BGP route reflector [[RFC4456](#)]. While the proposed modification cannot be used for inter-domain coordination of traffic filtering, it greatly simplifies distribution of intra-domain traffic filtering policies in an autonomous system with a large number of border routers having complex BGP policies. By relaxing the validation procedure for IBGP, the proposed modification allows flow specifications to be distributed in a standard and scalable manner throughout an autonomous system.

[4.](#) Revised Validation Procedure

Step (a) of the validation procedure specified in [RFC 5575, section 6](#) is redefined as follows:

- a) One of the following conditions MUST hold true:
 - o The originator of the flow specification matches the originator of the best-match unicast route for the destination prefix embedded in the flow specification.
 - o The AS_PATH and AS4_PATH attribute of the flow specification are empty.
 - o The AS_PATH and AS4_PATH attribute of the flow specification does not contain AS_SET and AS_SEQUENCE segments.

An empty AS_PATH and AS4_PATH attribute indicates per [[RFC4271](#)] that the flow specification NLRI originated in the same autonomous system as the local BGP speaker. Similarly, lack of AS_SET and AS_SEQUENCE segments within an AS_PATH and AS4_PATH attribute that is not empty indicates that the flow specification NLRI originated in the same autonomous system as the local BGP speaker but that the autonomous system includes a BGP confederation [[RFC5065](#)]. With this proposed modification to the [RFC 5575](#) validation procedure, it is now possible for an IBGP peer that is not within the data forwarding path to originate flow specification NLRIs. This applies with and without the presence of a BGP confederation within the autonomous system.

Further, [RFC 5575](#) states that "BGP (flow specification) implementations MUST also enforce that AS_PATH attribute of a route received via the External Border Gateway Protocol (eBGP) contains the neighboring AS in the left-most position of the AS_PATH attribute". This rule is not valid for all topologies. For example, it prevents exchange of BGP flow specification NLRIs at Internet exchanges with BGP route servers. Therefore, this document also redefines the [RFC 5575](#) AS_PATH and AS4_PATH validation procedure referenced above as follows.

BGP flow specification implementations MUST enforce that the last AS added within the AS_PATH and AS4_PATH attribute of a EBGP learned flow specification NLRI MUST match the last AS added within the AS_PATH and AS4_PATH attribute of the best-match unicast route for the destination prefix embedded in the flow specification. This proposed modification enables the exchange of BGP flow specification NLRIs at Internet exchanges with BGP route servers while at the same time, for security reasons, prevents an EBGP peer from advertising an inter-domain flow specification for a destination prefix that it does not provide reachability information for. Note, comparing only the

last ASNs is sufficient for EBGP learned flow specification NLRIs. Requiring a full AS_PATH and AS4_PATH match would limit origination of inter-domain flow specifications to the origin (or first) AS of the best-match unicast route for the destination prefix embedded in the flow specification only. As such, a full AS_PATH and AS4_PATH validity check may prevent transit ASNs from originating inter-domain flow specifications which is not desirable.

5. Security Considerations

No new security issues are introduced by relaxing the validation procedure for IBGP learned flow specifications. With this proposal, the security characteristics of BGP flow specifications remain equivalent to the existing security properties of BGP unicast routing. Traffic flow specifications learned from IBGP peers are trusted, hence, its not required to validate that the originator of an intra-domain traffic flow specification matches the originator of the best-match unicast route for the flow destination prefix. Conversely, this proposal continues to enforce the validation procedure for EBGP learned traffic flow specifications. In this way, the security properties of [RFC 5575](#) are maintained such that an EBGP peer cannot cause a denial-of-service attack by advertising an inter-domain flow specification for a destination prefix that it does not provide reachability information for.

6. IANA Considerations

This document has no actions for IANA.

7. Normative References

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

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

[RFC4456] Bates, T., Chen, E., and Chandra, R., "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", [RFC 4456](#), April 2006.

[RFC5065] Traina, P., McPherson, D., and Scudder, J., "Autonomous System Confederations for BGP", August 2007.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J.,

8. Acknowledgements

The authors would like to thank Han Nguyen for his direction on this work as well as Waqas Alam, Keyur Patel, Robert Raszuk, Eric Rosen and Shyam Sethuram for their review comments.

9. Authors' Addresses

James Uttaro
AT&T
200 S. Laurel Avenue
Middletown, NJ 07748
USA

Email: ju1738@att.com

Clarence Filsfils
Cisco
Brussels 1000
BE

Email: cf@cisco.com

Pradosh Mohapatra
Cisco
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: pmohapat@cisco.com

David J. Smith
Cisco
111 Wood Avenue South
Iselin, NJ 08830
USA

E-mail: djsmith@cisco.com

