

Internet Engineering Task Force
Internet-Draft
Updates: [5575](#) (if approved)
Intended status: Standards Track
Expires: January 15, 2014

H. Jeng
AT&T
J. Haas
Y. Rekhter
J. Zhang
Juniper Networks
July 14, 2013

Multicast Distribution Control Signaling draft-rekhter-mdcs-00

Abstract

This document describes a mechanism whereby the BGP Flow Specification NLRI format may be utilized to distribute multicast Control Plane filters. This mechanism is called Multicast Distribution Control Signaling (MDCS).

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 15, 2014.

Copyright Notice

Copyright (c) 2013 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.	Introduction	3
2.	Specification of Requirements	3
2.1.	Multicast Distribution Control Signaling	3
2.2.	An example of configuration on ERs	6
3.	Summary of Updates to BGP Flowspec	7
4.	IANA Considerations	7
5.	Security Considerations	7
6.	Acknowledgements	7
7.	References	7
7.1.	Normative References	7
7.2.	Informative References	8
	Authors' Addresses	8

1. Introduction

Consider a content provider that wants to deliver a particular content to a set of customers/subscribers, where the provider and the subscribers are connected by an IP service provider and the content is distributed using multicast connectivity. The content provider may wish to restrict delivery of the content to a subset of the subscribers in a centralized fashion.

For the purpose of this document we assume that a content provider consists of one or more Content Servers, and one or more Content Distribution Controllers. While this document assumes communication between Content Servers and Content Distribution Controllers, the procedures for implementing such communication is outside the scope of this document.

Content Servers are connected to one or more IP service providers (ISPs) that are offering multicast delivery of the content to the subscribers of the content provider. Content providers use these ISPs to deliver content to their subscribers.

Subscribers are connected to the Edge Routers (ERs) of the ISP. Note that the multicast connectivity service provided by the ISP extends all the way to the ERs. Such service could be provided by either deploying IP multicast natively, or with some tunneling mechanism like AMT, or by a combination of both within the ISP. However, between the ERs and the subscribers there may, or may not be multicast connectivity.

For further information, see [[geo-dist](#)].

2. 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 [RFC 2119](#) [[RFC2119](#)].

2.1. Multicast Distribution Control Signaling

Multicast distribution control signaling is intended to enforce exclusion/inclusion policies of a content provider, and specifically to prevent a subscriber from accessing a particular multicast channel carrying a particular content provided by the content provider if the subscriber obtained the information about this channel through some illegitimate means.

Multicast distribution control signaling for a particular content is

originated by Content Distribution Controller(s), and uses BGP Flow Spec [[RFC5575](#)] as follows:

For a particular content carried over a particular (S, G) multicast flow the Content Distribution Controller responsible for that content originates a BGP Flow Spec route. This route is carried using BGP multi-protocol capabilities [[RFC4760](#)] with AFI 1 (for IPv4) or 2 (for IPv6), and the MCAST-FLOWSPEC SAFI. The NLRI of the route carries S in the Source Prefix component (with length of 32 for IPv4 or 128 for IPv6), and G in the Destination Prefix component (with length of 32 for IPv4 or 128 for IPv6).

This route is ultimately propagated to the ER of the ISP connected to the content provider.

An ER that receives BGP Flow Spec routes carrying the multicast distribution control information applies it to PIM and/or IGMP messages the ER receives from the subscribers connected to that ER. (Note that such IGMP messages may be encapsulated in MDT messages.) Specifically, the ER, based on the information received in the BGP Flow Spec routes, decides whether to accept (or reject) a particular PIM or IGMP Join received on one of its subscriber's ports, as follows:

As a Content Distribution Controller originates a BGP Flow Spec route for a particular (S, G) multicast flow, such a route will carry one or more Route Targets [[RFC4360](#)], which will ultimately control inclusion/exclusion of that flow on individual ports of ERs that receive this route.

Each subscriber port on an ER is associated with one or more zones. For each zone that a port belongs to, the port is provisioned with two sets of RTs associated with that zone - the inclusion set is for allowing to accept PIM or IGMP Join for some content (or to be more precise for the (S, G) flow that carries that content), and the exclusion set is for disallowing to accept PIM or IGMP Joins for some other content. All those RTs (of all subscribers ports) control import of BGP Flow Spec routes by the ER.

Note that the RTs associated with the subscriber port are ordered. This permits configurations that accommodate include or exclude policies of zones of differing geographic size or overlap. See below for an example.

If the RTs carried by a given BGP Flow Spec route carrying multicast distribution control signaling match the inclusion set of RTs associated with a given port on an ER, then PIM or IGMP Joins for the (S, G) carried in the route and received from the subscriber(s)

connected to that port SHOULD be accepted by the ER. If the RTs carried by the route match the exclusion set, then PIM or IGMP Joins for the (S, G) carried in the route MUST NOT be accepted when received from the subscriber(s) connected to that port. (See example section below.)

Each subscriber port on an ER is provisioned with the default inclusion/exclusion policy that controls acceptance (or rejection) of PIM or IGMP Join messages in the absence of any multicast distribution control signaling. In the former case, in the absence of any multicast distribution signaling, subscribers connected to that port may receive any multicast flow. In the latter case, in the absence of any multicast distribution control signaling, subscribers connected to that port may receive no multicast flows. BGP Flow Spec routes that carry multicast distribution control signaling modify such default behavior.

Once a Content Distribution Controller determines that a particular (S, G) multicast stream no longer used to carry a particular content, the Content Distribution Controller withdraws the BGP Flow Spec route that carries multicast distribution control information for that content.

Note that while [[RFC5575](#)] uses the information carried in BGP Flow Spec routes for the purpose of Data Plane filtering, this document uses this information for the purpose of filtering multicast Control Plane traffic (PIM or IGMP).

To constrain the distribution of BGP Flow Spec routes that carry multicast distribution control information to only the relevant ERs, the ERs MAY originate Route Target Constraint (RTC) routes that carry the RTs that control import of the BGP Flow Spec routes on these ERs.

To constrain the import of these RTC routes to only the Content Distribution Controllers, the Content Distribution Controllers are configured with one or more RTs. These RTs control import by the Content Distribution Controller(s) of the RTC routes originated by the ERs. Furthermore, the Content Distribution Controllers MAY themselves originate RTC routes that carry the import RT(s) configured on these Content Distribution Controllers, and that control import of RTC routes by these Content Distribution Controllers.

This document assumes that if a given content provider has multiple Content Distribution Controllers, then all of these Controllers are provisioned with the same RT(s) that control import of the RTC routes originated by the ERs. Furthermore, this document assumes that if a given ISP is providing (multicast) connectivity service to more than

one content provider, then the RTC routes originated by any of the ERs of that ISP MUST carry the set union of the import RTs used by the Content Distribution Controllers of all of these content providers.

RTs carried by routes with AFI 1 and MCAST-FLOWSPEC SAFI SHOULD NOT be re-used by routes with any other AFI and/or SAFI. Likewise, RTs carried by routes with AFI 2 and MCAST-FLOWSPEC SAFI SHOULD NOT be re-used by routes with any other AFI and/or SAFI. Furthermore, RTs carried by routes with AFI 1 and SAFI 132 (AFI/SAFI used by RTC routes) SHOULD NOT be re-used by routes with any other AFI and/or SAFI.

Note that while [\[RFC4684\]](#) uses RTC routes to constrain distribution of VPN-IP routes [\[RFC4364\]](#), this document uses RTC routes to constrain distribution of BGP Flow Spec routes, and also to (recursively) constrain distribution of RTC routes themselves.

2.2. An example of configuration on ERs

Consider an ER in Manhattan that has a port that is provisioned with the following import RTs:

```
<include-manhattan, exclude-manhattan, include-nyc, exclude-  
nyc, include-east, exclude-east, include-usa, exclude-usa>
```

When the ER receives a Flow Spec route with <exclude-nyc, include-manhattan, include-usa> RTs, the ER first try to match "include-manhattan" or "exclude-manhattan" (the first ones on the list) - and the result is "include-manhattan". Therefore, the (S, G) carried in the Flow Spec route is allowed on that port of the ER.

Consider another ER in Boston that has a port that is provisioned with the following import RTs:

```
<include-cambridge, exclude-cambridge, include-bos, exclude-  
bos, include-east, exclude-east, include-usa, exclude-usa>
```

The above mentioned Flow Spec route will be imported (due to the include-usa RT), and will result in the (S, G) carried in the flow Spec route to be allowed on that port of the ER.

Now consider a different Flow Spec route with the <exclude-usa, include-bos, include-nyc, exclude-manhattan> RTs. The (S, G) carried in the route will be disallowed in Manhattan, allowed in Boston, and allowed in Queens (as the route will match the "include-nyc" RT).

3. Summary of Updates to BGP Flowspec

As described above, this document makes small changes to the BGP Flow Specification mechanism when carried using the MCAST-FLOWSPEC SAFI:

- o Destination addresses will contain a multicast group rather than a unicast destination.
- o Flow specification routes for this SAFI are used for filtering multicast Control Plane traffic rather than the matching multicast traffic itself.
- o Flow specification routes for this SAFI will carry one or more Route Target extended communities.
- o Flow specification component types not applicable to signaling multicast Control Plane traffic MUST be ignored. E.g.: ICMP type, ICMP code, TCP flags, Fragment.

4. IANA Considerations

This document defines a new BGP Subsequent Address Family Identifier (SAFI) value, MCAST-FLOWSPEC. The authors request assignment of a value from the First Come, First Served portion of this registry.

5. Security Considerations

TBD

6. Acknowledgements

The authors would like to thank Han Nguyen for his contributions to this document.

7. References

7.1. Normative References

- [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.

- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", [RFC 4684](#), November 2006.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](#), August 2009.

7.2. Informative References

- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [geo-dist] Jeng, H., Haas, J., Rekhter, Y., and J. Zhang, "Multicast Geo Distribution Control", [draft-rekhter-geo-distribution-control-03.txt](#) (work in progress), 2013.

Authors' Addresses

Huajin Jeng
AT&T

Email: hj2387@att.com

Jeffrey Haas
Juniper Networks
1194 N. Mathida Ave.
Sunnyvale, CA 94089
US

Email: jhaas@juniper.net

Yakov Rekhter
Juniper Networks
1194 N. Mathida Ave.
Sunnyvale, CA 94089
US

Email: yakov@juniper.net

Jeffrey (Zhaohui) Zhang
Juniper Networks
1194 N. Mathida Ave.
Sunnyvale, CA 94089
US

Email: zzhang@juniper.net

