IDR Working Group
Internet-Draft
Updates: 4684 (if approved)
Intended status: Standards Track
Expires: May 16, 2016

E. Rosen, Ed.
Juniper Networks, Inc.
K. Patel
Cisco Systems, Inc.
J. Haas
Juniper Networks, Inc.
R. Raszuk
Bloomberg LP

November 13, 2015

Route Target Constrained Distribution of Routes with no Route Targets
draft-ietf-idr-rtc-no-rt-04.txt

Abstract

There are a variety of BGP-enabled services in which the originator of a BGP route may attach one or more "Route Targets" to the route. By means of a procedure known as "RT Constrained Distribution" (RTC), a given BGP speaker (call it "B") can announce the set of RTs in which it has interest. The implication is that if a particular route (call it "R") carries any RTs at all, BGP speaker B wants to receive route R if and only if B has announced interest in one of the RTs carried by R. However, if route R does not carry any RTs at all, prior specifications do not make it clear whether B's use of RTC implies that it does not want to receive route R. This has caused interoperability problems in the field, as some implementations of RTC do not allow B to receive R, but some services presuppose that B will receive R. This document updates RFC 4684 by clarifying the effect of the RTC mechanism on routes that do not have any RTs.

Table of Contents

1. Introduction

A BGP route can carry a particular type of BGP path attribute known as an "Extended Communities Attribute" [RFC4360]. Each such attribute can contain a variable number of typed communities. Certain typed communities are known as "Route Targets" (RTs) ([RFC4360], [RFC4364]).

[RFC4684] defines a procedure, known as "RT Constrained Distribution" (RTC) that allows a BGP speaker to advertise its interest in a particular set of RTs. It does so by advertising "RT membership information". (See [RFC4684] for details.) It may advertise RT membership for any number of RTs. By advertising membership for a particular RT, a BGP speaker declares that it is interested in receiving BGP routes that carry that RT.

If RTC is enabled on a particular BGP session, the session must be provisioned with the set of "address family" and "subsequent address family" values (AFI/SAFIs) to which RTC is to be applied. In [RFC4684] it is implicitly assumed that RTC will only be applied to AFI/SAFIs for which all the routes carry RTs. When this assumption

is true, the RTC semantics are clear. A BGP speaker advertising its interest in RT1, RT2, ..., RTk is saying that, for the AFI/SAFIs to which RTC is being applied, it is interested in any route that carries at least one of those RTs, and it is not interested in any route that does not carry at least one of those RTs.

However, [RFC4684] does not specify how the RTC procedures are to be applied to AFI/SAFIs whose routes sometimes carry RTs and sometimes do not. Consider a BGP session between routers R1 and R2, where R1 has advertised its interest in RT1, RT2, ..., RTk, and RTC is being applied to a particular AFI/SAFI. Suppose R2 has a route of that AFI/SAFI, and that route carries no RTs. Should R2 advertise this route to R1 or not?

There are two possible answers to this question, each of which seems prima facie reasonable:

- No, R2 should not advertise the route, because it belongs to an AFI/SAFI to which RTC is being applied, and the route does not carry any of the RTs in which R1 is interested.

- Yes, R2 should advertise the route; since the route carries no RTs, the intention of the route's originator is that the distribution of the route not be constrained by the RTC mechanism.

As might be expected, "one size does not fit all". The best answer depends upon the particular deployment scenario, and upon the particular AFI/SAFI to which RTC is being applied.

Section 3 defines a default behavior for existing AFI/SAFIs. This default behavior ensures proper operation when RTC is applied to an existing AFI/SAFI. The default behavior may of course be overridden by local policy.

Section 3 also defines a default "default behavior" for new AFI/SAFIs. When a new AFI/SAFI is defined, the specification defining it may specify a different default behavior; otherwise the default default behavior will apply.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Some Deployment Scenarios

The lack of a clearly defined default behavior for applying RTC to routes that carry no RTs is problematic in at least three scenarios.

- [RFC6037] describes a deployed Multicast VPN (MVPN) solution. It defines a BGP SAFI known as "MDT-SAFI". Routes with this SAFI may carry RTs, but are not required to do so. In order for the procedures of [RFC6037] to work properly, if an MDT-SAFI route does not carry any RTs, the distribution of that route MUST NOT be constrained by RTC. However, if an MDT-SAFI route does carry one or more RTs, its distribution SHOULD be constrained by RTC.

- [GTM] specifies a way to provide "Global Table Multicast" (as opposed to VPN multicast), using procedures that are very similar to those described in [RFC6513] and [RFC6514] for MVPN. In particular, it uses routes of the MCAST-VPN SAFI that is defined in [RFC6514]. When used for MVPN, each MCAST-VPN route carries at least one RT. However, when used for Global Table Multicast, it is optional for certain MCAST-VPN routes to carry RTs. In order for the procedures of [GTM] to work properly, if an MCAST-VPN route does not carry any RTs, the distribution of that route MUST NOT be constrained by RTC.

- Typically, Route Targets have been carried only by routes that are distributed as part of a VPN service (or the Global Table Multicast service mentioned above). However, it may be desirable to be able to place RTs on non-VPN routes (e.g., on unicast IPv4 or IPv6 routes) and then to use RTC to constrain the delivery of the non-VPN routes. For example, if a BGP speaker desires to receive only a small set of IPv4 unicast routes, and the desired routes carry one or more RTs, the BGP speaker could use RTC to advertise its interest in one or more of those RTs. In this application, the intention would be that any IPv4 unicast route not carrying an RT would be filtered. Note that this is the opposite of the behavior needed for the other use cases discussed in this section.

3. Default Behavior

In order to handle the use cases discussed in Section 2, this document specifies a default behavior for the case where RTC is applied to a particular AFI/SAFI, and some (or all) routes of that address family do not carry any RTs.

When RTC is applied, on a particular BGP session, to routes of the MDT-SAFI address family (SAFI=66, [RFC6037]), the default behavior

MUST be that routes that do not carry any RTs are distributed on that session.

When RTC is applied, on a particular BGP session, to routes of the MCAST-VPN address family (SAFI=5, [RFC6514], [GTM]), the default behavior MUST be that routes that do not carry any RTs are distributed on that session.

When RTC is applied, on a particular BGP session, to routes of other address families, the default behavior MUST be that routes without any RTs are not distributed on that session. This default "default behavior" applies to all AFI/SAFIs for which a different default behavior has not been defined.

A BGP speaker MAY be provisioned to apply a non-default behavior to a given AFI/SAFI. This is a matter of local policy.

4. IANA Considerations

This document contains no actions for IANA.

5. Security Considerations

The security considerations of [RFC4684] apply.

The procedures of this document may allow the distribution of certain SAFI-5 and SAFI-66 routes, in situations where some implementations of RTC would previously have prevented their distribution. However, it is necessary to distribute such routes in order for the applications using them to operate properly. Allowing the distribution of such routes does not create any new security considerations beyond those of the applications that use the routes.

6. References

6.1. Normative References

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

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <http://www.rfc-editor.org/info/rfc4360>.

[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, DOI 10.17487/RFC4684, November 2006, <http://www.rfc-editor.org/info/rfc4684>.

6.2. Informative References

[GTM] Zhang, J., Giulano, L., Rosen, E., Subramanian, K., and D. Pacella, "Global Table Multicast with BGP-MVPN Procedures", internet-draft draft-ietf-bess-mvpn-global-table-mcast-03, September 2015.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <http://www.rfc-editor.org/info/rfc4364>.

[RFC6037] Rosen, E., Ed., Cai, Y., Ed., and IJ. Wijnands, "Cisco Systems' Solution for Multicast in BGP/MPLS IP VPNs", RFC 6037, DOI 10.17487/RFC6037, October 2010, <http://www.rfc-editor.org/info/rfc6037>.

[RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <http://www.rfc-editor.org/info/rfc6513>.

[RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <http://www.rfc-editor.org/info/rfc6514>.

Authors' Addresses

Eric C. Rosen (editor)
Juniper Networks, Inc.
10 Technology Park Drive
Westford, Massachusetts 01886
United States

Email: erosen@juniper.net

Keyur Patel
Cisco Systems, Inc.
170 Tasman Drive
San Jose, California 95134
United States

Email: keyupate@cisco.com

Jeffrey Haas
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, California 94089
United States

Email: jhaas@juniper.net

Robert Raszuk
Bloomberg LP
731 Lexington Ave
New York City, NY 10022
United States

Email: robert@raszuk.net