

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 9, 2015

E. Rosen, Ed.
Juniper Networks, Inc.
K. Patel
Cisco Systems, Inc.
J. Haas
Juniper Networks, Inc.
R. Raszuk
Mirantis Inc.
January 5, 2015

Route Target Constrained Distribution of Routes with no Route Targets
draft-ietf-idr-rtc-no-rt-00.txt

Abstract

BGP routes sometimes carry an "Extended Communities" path attribute. An Extended Communities path attribute can contain one or more "Route Targets" (RTs). By means of a procedure known as "RT Constrained Distribution" (RTC), a BGP speaker can send BGP UPDATE messages that express its interest in a particular set of RTs. Generally, RTC has been applied only to address families whose routes always carry RTs. When RTC is applied to such an address family, a BGP speaker expressing its interest in a particular set of RTs is indicating that it wants to receive all and only the routes of that address family that have at least one of the RTs of interest. However, there are scenarios in which the originator of a route chooses not to include any RTs at all, assuming that the distribution of a route with no RTs at all will be unaffected by RTC. This has led to interoperability problems in the field, where the originator of a route assumes that RTC will not affect the distribution of the route, but intermediate BGP speakers refuse to distribute that route because it does not carry any RT of interest. The purpose of this document is to clarify the effect of the RTC mechanism on routes that do not have any RTs.

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 July 9, 2015.

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 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.	Some Deployment Scenarios	4
3.	Default Behavior	4
4.	IANA Considerations	5
5.	Security Considerations	5
6.	References	5
6.1.	Normative References	5
6.2.	Informative References	5
	Authors' Addresses	6

[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" (AFI/SAFIs) values to which RTC is to be applied. In [\[RFC4684\]](#) it is implicitly assumed that RTC will only be applied to AFI/SAFIs where 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 address families 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 different answers to this question, each of which seems *prima facie* reasonable:

- o No, R2 should not advertise the route, because it belongs to an AFI/SAFI to which RTC is being applied, and the route does carry any of the RTs in which R1 is interested.
- o 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", and 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 each existing AFI/SAFI. This default behavior will ensure proper operation of that AFI/SAFI when RTC is applied. The default behavior may of course be overridden by a 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" are to be interpreted as described in [\[RFC2119\]](#).

2. Some Deployment Scenarios

There are at least three deployment scenarios where lack of a clearly defined default behavior for RTC is problematic.

- o [\[RFC6037\]](#) describes a deployed Multicast VPN (MVPN) solution. It defines a BGP address family known as "MDT-SAFI". Routes of this address family may carry RTs, but are not required to do so. In order for the [RFC6037](#) procedures 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 may be constrained by RTC.
- o [\[GTM\]](#) specifies a way to provide "global table" (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 address family 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 route types 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.
- o Typically, Route Targets have been carried only by routes that are distributed as part of a VPN service. 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 3](#), this document specifies a default behavior for the case where RTC is applied to a particular address family (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), the default behavior is 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), the default behavior is 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 is 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

No security considerations are raised by this document beyond those already discussed in [[RFC4684](#)].

6. References

6.1. 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.

6.2. Informative References

- [GTM] Zhang, J., Giulano, L., Rosen, E., Subramanian, K., Pacella, D., and J. Schiller, "Global Table Multicast with BGP-MVPN Procedures", internet-draft [draft-ietf-l3vpn-mvpn-global-table-mcast-00](#), July 2014.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.

- [RFC6037] Rosen, E., Cai, Y., and IJ. Wijnands, "Cisco Systems' Solution for Multicast in BGP/MPLS IP VPNs", [RFC 6037](#), October 2010.
- [RFC6513] Rosen, E. and R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", [RFC 6513](#), February 2012.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", [RFC 6514](#), February 2012.

Authors' Addresses

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

Email: erosen@juniper.net

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

Email: keyupate@cisco.com

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

Email: jhaas@juniper.net

Robert Raszuk
Mirantis Inc.
615 National Ave. #100
Mountain View, California 94043
US

Email: robert@raszuk.net

