

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 14, 2012

K. Patel
Cisco Systems
R. Raszuk
NTT MCL Inc.
M. Djernaes
Juniper Networks
J. Dong
M. Chen
Huawei Technologies
June 12, 2012

IPv6 Extensions for Route Target Distribution
draft-ietf-idr-bgp-ipv6-rt-constrain-02

Abstract

The current route target distribution specification described in [RFC4684](#) defines Route Target NLRIs of maximum length of 12 bytes. The IPv6 specific Route Target extended community is defined in [\[RFC5701\]](#) as length of 20 bytes. Since the current specification only supports prefixes of maximum length of 12 bytes, the lack of an IPv6 specific Route Target reachability information may be a problem when an operator wants to use this application in a pure IPv6 environment. This document defines an extension that allows BGP to exchange longer length IPv6 Route Target prefixes.

Requirements Language

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](#)].

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 December 14, 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.	Introduction	4
2.	BGP IPv6 Constrained Route Target Capability	4
3.	IPv6 Constrained Route Target NLRI Advertisements	4
4.	IANA Considerations	5
5.	Security Considerations	5
6.	Acknowledgements	5
7.	References	5
7.1.	Normative References	5
7.2.	Informative References	6
	Authors' Addresses	6

1. Introduction

The current constrained route distribution specification defined in [RFC4684] supports prefixes with a maximum length of 12 bytes. The prefix length needs to be extended to support the IPv6 specific Route Target extended community defined in [RFC5701] which is 20 bytes in length. This document defines an extension to the current constrained route distribution specification that allows BGP speakers to distribute longer length Route Target prefixes. A new BGP capability known as BGP IPv6 Constrained Route Target capability is defined as part of extension that allows an exchange of longer length Route Target prefixes. BGP speakers that do not exchange this capability MUST use Route Target NLRIs of maximum length of 12 bytes. In this way, the current extension would preserve the backward compatibility with [RFC4684].

2. BGP IPv6 Constrained Route Target Capability

The "BGP IPV6 Constrained Route Target Capability" is a new BGP capability [RFC5492]. The Capability code for this capability is specified in the IANA Considerations section of this document. The Capability length field of this capability is zero.

By advertising this capability to a peer, a BGP speaker conveys to the peer that the speaker support the longer length Route Target prefixes and the related procedures described in this document.

3. IPv6 Constrained Route Target NLRI Advertisements

Route Target membership NLRI is advertised in BGP UPDATE messages using the MP_REACH_NLRI and MP_UNREACH_NLRI attributes as defined in [RFC4760]. The NLRI field in the MP_REACH_NLRI and MP_UNREACH_NLRI is a prefix of 0 to 24 octets, encoded as defined in [Section 4 of \[RFC4760\]](#) for all the constrained route distribution.

This prefix is structured as follows:

```

+-----+
| origin as      (4 octets) |
+-----+
| route target  (8 or 20 octets)|
~                               ~
|                               |
+-----+

```

Except for the default route target, which is encoded as a zero-length prefix, the minimum prefix length is 32 bits. As the

origin-AS field cannot be interpreted as a prefix.

Route targets can then be expressed as prefixes, where, for instance, a prefix would encompass all route target extended communities assigned by a given Global Administrator [[RFC4360](#)] and [[RFC5701](#)]. Alternatively, route target prefixes could be aggregated however if done so, then only the Local Administrator field of the Route Target can be aggregated. Route Target Type and the Global Administrator Route Target fields MUST not be aggregated.

The default route target can be used to indicate to a peer the willingness to receive all VPN route advertisements such as, for instance, the case of a route reflector speaking to one of its PE router clients.

4. IANA Considerations

This document defined the IPV6 Constrained Route Target Capability for BGP. The Capability code needs to be assigned by the IANA.

5. Security Considerations

This extension to [[RFC4684](#)] does not change the underlying security issues inherent in the existing BGP and [[RFC4684](#)].

6. Acknowledgements

The authors would like to thank Pedro Marques, John Scudder, Alton Lo and Zhenqiang Li for discussions and review.

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

[RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), February 2009.

[RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", [RFC 5701](#), November 2009.

7.2. Informative References

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.

Authors' Addresses

Keyur Patel
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: keyupate@cisco.com

Robert Raszuk
NTT MCL Inc.
101 S Ellsworth Avenue Suite 350
San Mateo, CA 94401
USA

Email: robert@raszuk.net

Martin Djernaes
Juniper Networks
1194 N. Mathilda Avenue
Sunnyvale, CA 94089
USA

Email: mdjernaes@juniper.net

Jie Dong
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

Mach(Guoyi) Chen
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing 100095
China

Email: mach.chen@huawei.com

