

Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: December 27, 2009

J. Scudder  
Juniper Networks  
J. Uttaro  
AT&T  
P. Mohapatra  
Cisco Systems  
June 25, 2009

RT-Constrain Lite for Provider Edge Routers  
draft-scudder-idr-rt-constrain-lite-00.txt

## Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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/1id-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 27, 2009.

## Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the

Internet-Draft

RT-Constrain Lite

June 2009

document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

## Abstract

[RFC 4684](#), "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)" provides a powerful and general means for BGP speakers to exchange and propagate Route Target reachability information which is used for cooperative route filtering. However, the complexity of implementing the entire specification may have impeded its widespread deployment. This document specifies the subset of functionality which is required for a provider edge router ("PE") to originate Route Target NLRI. Such PEs need not implement any filtering functionality.

## 1. Introduction

[[RFC4684](#)], "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)" provides a powerful and general means for BGP speakers to exchange and propagate Route Target reachability information which is used for cooperative route filtering. However, the complexity of implementing the entire specification may have impeded its widespread deployment. We observe that the functionality required for a BGP speaker functioning solely as a provider edge router ("PE") is substantially simpler than that required for a speaker which serves as a route reflector or ASBR. Specifically, the PE need only implement the ability to send Route Target Membership NLRI; it need not have the ability to receive, store and filter upon such information.

This document specifies the subset of functionality which is required for a PE to originate Route Target NLRI. Since this document simply specifies a subset, any BGP implementation which conforms with [[RFC4684](#)] by definition also conforms with this specification.

## [1.1.](#) 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](#)].

## [2.](#) Route Target Membership NLRI Advertisements

A PE implementing this specification advertises its Route Targets of interest as Route Target membership NLRI in BGP UPDATE messages using the MP\_REACH\_NLRI and MP\_UNREACH\_NLRI attributes [[RFC4760](#)]. The [AFI, SAFI] value pair used to identify this NLRI is (AFI=1, SAFI=132).

The NLRI field in the MP\_REACH\_NLRI and MP\_UNREACH\_NLRI is a prefix of 0 to 96 bits, encoded as defined in Sections [3](#) and [4](#) of [[RFC4760](#)].

This prefix is structured as follows:

```
+-----+
| origin AS      (4 octets) |
+-----+
| route target   (8 octets) |
+-----+
|                     |
+-----+
```

Except for the default route target, which is encoded as a zero-length prefix, the minimum prefix length is 32 bits, since the origin AS field cannot be interpreted as a prefix.

Although route targets can be expressed as prefixes as discussed in [[RFC4684](#)], this specification does not mandate that such functionality be provided. An implementation MAY choose to implement the ability to advertise route target prefixes.

Although the default route target can be used to indicate to a peer the willingness to receive all VPN route advertisements as discussed in [[RFC4684](#)], this functionality is of dubious utility for a PE and thus this specification does not mandate that such functionality be provided. An implementation MAY choose to implement the ability to advertise the default route target.

### [3.](#) Capability Advertisement

A BGP speaker that wishes to advertise Route Target membership information **MUST** use the Multiprotocol Extensions Capability [[RFC5492](#)] Code, as defined in [[RFC4760](#)], to advertise the corresponding (AFI, SAFI) pair, and **MUST NOT** advertise Route Target membership information unless its peer has similarly advertised the (AFI, SAFI) pair.

### [4.](#) Operation

A BGP speaker implementing this specification **MAY** ignore all received Route Target Membership NLRI routes. Such routes need not be stored, they **MAY** be completely discarded without further processing. A consequence of this is that a BGP speaker implementing this specification **MAY** advertise its VPN NLRI without regard to what Route Target membership information its peers may or may not have advertised.

A BGP speaker implementing this specification **MUST** advertise a Route Target Membership NLRI for each Route Target which it has been configured to import into a local VRF. When the speaker's configuration is updated to add or remove a Route Target import, the speaker **MUST** generate Route Target Membership NLRI updates (advertisements and/or withdrawals) to convey the necessary changes.

As a hint that initial RT membership exchange is complete, implementations **SHOULD** generate an End-of-RIB marker, as defined in [[RFC4724](#)], for the Route Target membership (afi, safi), regardless of whether graceful restart is enabled on the BGP session.

### [5.](#) Deployment Observations

We observe that when a BGP speaker supporting [[RFC4684](#)] and acting as a route reflector ("the RR") peers with a BGP speaker which implements this specification ("the PE"), the PE can be expected to send all its VPN routes to the RR just as it would if no Cooperative

Route Filtering were in use. The RR would not be expected to apply any filtering to those routes as a consequence of this specification or [[RFC4684](#)]. However, the RR would be expected to build an outbound filter towards the PE based on Route Target membership information received from the PE. This is a consequence of the normal operation of [[RFC4684](#)]; refer to that specification for more detail.

## [6.](#) Acknowledgements

This document relies entirely on the functionality defined in [[RFC4684](#)]. As such, thanks are due to the authors of that document. Jim Guichard also made valuable contributions.

## [7.](#) IANA Considerations

This memo includes no request to IANA.

Scudder, et al.

Expires December 27, 2009

[Page 4]

---

Internet-Draft

RT-Constrain Lite

June 2009

## [8.](#) Security Considerations

This specification makes no changes to the security considerations of [[RFC4684](#)].

## [9.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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.
- [RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", [RFC 4724](#), January 2007.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter,

"Multiprotocol Extensions for BGP-4", [RFC 4760](#),  
January 2007.

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

#### Authors' Addresses

John G. Scudder  
Juniper Networks

Email: [jgs@juniper.net](mailto:jgs@juniper.net)

Jim Uttaro  
AT&T

Email: [uttaro@att.com](mailto:uttaro@att.com)

Pradosh Mohapatra  
Cisco Systems

Email: [pmohapat@cisco.com](mailto:pmohapat@cisco.com)