

Route Refresh Capability for BGP-4

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This document defines a new Border Gateway Protocol (BGP) capability termed 'Route Refresh Capability', which would allow the dynamic exchange of route refresh request between BGP speakers and subsequent re-advertisement of the respective Adj-RIB-Out. One possible application of this capability is to facilitate non-disruptive routing policy changes.

1. Introduction

Currently there does not exist a mechanism in BGP-4 [BGP-4] to dynamically request a re-advertisement of the Adj-RIB-Out from a BGP peer. When the inbound routing policy for a peer changes, all prefixes from that peer must be somehow made available and then re-examined against the new policy. To accomplish this, a commonly used approach, known as 'soft-reconfiguration', is to store an unmodified copy of all routes from that peer at all times, even though routing policies do not change frequently (typically no more than a couple times a day). Additional memory and CPU are required to maintain these routes.

This document proposes an alternative solution that avoids the additional maintenance cost. More specifically, it defines a new BGP capability termed 'Route Refresh Capability', which would allow the dynamic exchange of route refresh request between BGP speakers and subsequent re-advertisement of the respective Adj-RIB-Out.

2. Route Refresh Capability

To advertise the Route Refresh Capability to a peer, a BGP speaker uses BGP Capabilities Advertisement [BGP-CAP]. This capability is advertised using the Capability code 2 and Capability length 0.

By advertising the Route Refresh Capability to a peer, a BGP speaker conveys to the peer that the speaker is capable of receiving and properly handling the ROUTE-REFRESH message (as defined in Section 3) from the peer.

3. Route-REFRESH Message

The ROUTE-REFRESH message is a new BGP message type defined as follows:

Type: 5 - ROUTE-REFRESH

Message Format: One <AFI, SAFI> encoded as

| 0 | 7 | 15 | 23 | 31 |
|---------------------------|-----|----|------|------|
| +-----+-----+-----+-----+ | | | | |
| | AFI | | Res. | SAFI |
| +-----+-----+-----+-----+ | | | | |

The meaning, use and encoding of this <AFI, SAFI> field is the same as defined in [BGP-MP, sect. 7]. More specifically,

AFI - Address Family Identifier (16 bit).

Res. - Reserved (8 bit) field. Should be set to 0 by the sender and ignored by the receiver.

SAFI - Subsequent Address Family Identifier (8 bit).

4. Operation

A BGP speaker that is willing to receive the ROUTE-REFRESH message from its peer should advertise the Route Refresh Capability to the peer using BGP Capabilities advertisement [BGP-CAP].

A BGP speaker may send a ROUTE-REFRESH message to its peer only if it has received the Route Refresh Capability from its peer. The <AFI, SAFI> carried in such a message should be one of the <AFI, SAFI> that the peer has advertised to the speaker at the session establishment time via capability advertisement.

If a BGP speaker receives from its peer a ROUTE-REFRESH message with the <AFI, SAFI> that the speaker didn't advertise to the peer at the session establishment time via capability advertisement, the speaker shall ignore such a message. Otherwise, the BGP speaker shall re-advertise to that peer the Adj-RIB-Out of the <AFI, SAFI> carried in the message, based on its outbound route filtering policy.

5. Security Considerations

This extension to BGP does not change the underlying security issues.

6. Acknowledgments

The concept of Route Refresh proposed is similar to the one used in IDRP.

The author would like to thank Yakov Rekhter, Ravi Chandra, Srihari Ramachandra and Bruce Cole for their review and comments.

7. References

- [BGP-4] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.
- [BGP-MP] Bates, T., Chandra, R., Katz, D. and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 2858, June 2000.
- [BGP-CAP] Chandra, R. and J. Scudder, "Capabilities Advertisement with BGP-4", RFC 2842, May 2000.

8. Author's Address

Enke Chen
Redback Networks Inc.
350 Holger Way
San Jose, CA 95134

EMail: enke@redback.com

9. Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

