

Route Refresh Capability for BGP-4

[draft-ietf-idr-bgp-route-refresh-01.txt](#)

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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>.

2. Abstract

This document 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. One possible application of this capability is to facilitate non-disruptive routing policy changes.

3. 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.

4. 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 5](#)) from the peer.

5. 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. 8]. 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).

6. 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.

7. Security Considerations

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

8. 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.

9. 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 Rekhter, Y., 'Multiprotocol Extensions for BGP-4', work in progress

[BGP-CAP] Chandra, R., Scudder, J., 'Capabilities Advertisement with BGP-4', work in progress

10. Author Information

Enke Chen
Siara Systems Incorporated
1195 Borregas Ave.
Sunnyvale, Ca. 94089
e-mail: enkechen@siara.com

