

Network Working Group
Internet Draft
Intended Status: Standards Track
Expiration Date: June 17, 2012

K. Patel
E. Chen
B. Venkatachalapathy
Cisco Systems
December 16, 2011

Enhanced Route Refresh Capability for BGP-4
draft-ietf-idr-bgp-enhanced-route-refresh-01.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on June 17, 2012.

Copyright Notice

Copyright (c) 2011 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.

Abstract

In this document we enhance the existing BGP route refresh mechanisms to provide for the demarcation of the beginning and the ending of a route refresh. The enhancement can be used to facilitate on-line, non-disruptive consistency validations of BGP routing updates.

1. Introduction

It is sometimes necessary to perform routing consistency validations such as checking for possible missing withdraws between BGP speakers [[RFC4271](#)]. Currently such validations typically involve off-line, manual operations which can be tedious and time consuming.

In this document we enhance the existing BGP route refresh mechanisms [[RFC2918](#)] to provide for the demarcation of the beginning and the ending of a route refresh (which refers to the complete re-advertisement of the Adj-RIB-Out to a peer, subject to routing policies). The enhancement can be used to facilitate on-line, non-disruptive consistency validation of BGP routing updates.

1.1. Specification of Requirements

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 [[RFC2119](#)].

2. Protocol Extensions

The BGP protocol extensions introduced in this document include the definition of a new BGP capability, named "Enhanced Route Refresh Capability", and the specification of the message subtypes for the ROUTE-REFRESH message.

2.1. Enhanced Route Refresh Capability

The "Enhanced Route Refresh 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 supports the message subtypes for the ROUTE-REFRESH message and the related procedures described in this

document.

2.2. Subtypes for ROUTE-REFRESH Message

The "Reserved" field of the ROUTE-REFRESH message specified in [RFC2918] is re-defined as the "Message Subtype" with the following values:

- 0 - Normal route refresh request [RFC2918]
with/without ORF [RFC5291]
- 1 - Demarcation of the beginning of a route refresh
- 2 - Demarcation of the ending of a route refresh

The use of the message subtypes is described in the Operations section.

3. Operations

A BGP speaker that support the message subtypes for the ROUTE-REFRESH message and the related procedures SHOULD advertise the "Enhanced Route Refresh Capability".

The following procedures are applicable only if a BGP speaker has received the "Enhanced Route Refresh Capability" from a peer.

Before the speaker starts a route refresh that is either initiated locally, or in response to a "normal route refresh request" from the peer, the speaker MUST send a ROUTE-REFRESH message with the specified message subtype to mark the beginning of the route refresh. After the speaker completes the re-advertisement of the entire Adj-RIB-Out to the peer, it MUST send a ROUTE-REFRESH message with the specified message subtype to mark the ending of the route refresh.

Conceptually the "entire ADJ-RIB-Out" for a peer in this section refers to all the route entries in the "ADJ-RIB-Out" for the peer at the start of the route refresh. When a route entry in the "ADJ-RIB-Out" changes, the advertisement of the modified route entry (instead of the snapshot entry) would suffice.

In processing a ROUTE-REFRESH message from a peer, the BGP speaker MUST examine the "message subtype" field of the message and take the appropriate actions. The BGP speaker SHALL use the demarcations of the beginning and the ending of a route refresh to perform consistency validations of the updates received from the peer. All

the routes that were not re-advertised in the route refresh MUST be purged, and SHOULD be logged for further analysis.

4. Error Handling

This document defines a new NOTIFICATION error code:

Error Code	Symbolic Name
<TBD>	ROUTE-REFRESH Message Error

The following error subcodes are defined as well:

Subcode	Symbolic Name
1	Invalid Message Length

The error handling specified in this section is applicable only when a BGP speaker has received the "Enhanced Route Refresh Capability" from a peer.

When the BGP speaker detects an error while processing a ROUTE-REFRESH message with a non-zero "Message Subtype" field, it MUST send a NOTIFICATION message with Error Code "ROUTE-REFRESH Message Error". The Data field of the NOTIFICATION message MUST contain the complete ROUTE-REFRESH message.

If the length, excluding the fixed-size message header, of the ROUTE-REFRESH message with Message Subtype 1 and 2 is not 4, then the error subcode is set to "Invalid Message Length".

5. IANA Considerations

This document defines the Enhanced Route Refresh Capability for BGP. The Capability Code 70 has been assigned by the IANA.

In addition, this document defines an NOTIFICATION error code and several error subcodes for the ROUTE-REFRESH message. They need to be registered with the IANA.

6. Security Considerations

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

7. Acknowledgments

The authors would like to thank Pedro Marques, Pradosh Mohapatra, Robert Raszuk, Pranav Mehta, and Shyam Sethuram for discussions and review. The authors would like to thank Martin Djernaes, Jeff Haas, Ilya Varlashkin, Rob Shakir, Paul Jakma, Jie Dong, Qing Zeng, Albert Tian, and Jakob Heitz for their review and comments.

8. Normative References

- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC2918] Chen, E., "Route Refresh Capability for BGP-4", [RFC 2918](#), September 2000.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), February 2009.
- [RFC5291] Chen, E., and Rekhter, Y., "Outbound Route Filtering Capability for BGP-4", [RFC 5291](#), August 2008.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

9. Authors' Addresses

Keyur Patel
Cisco Systems

Email: keyupate@cisco.com

Enke Chen
Cisco Systems

Email: enkechen@cisco.com

Balaji Venkatachalapathy
Cisco Systems

Email: bvenkata@cisco.com