

Network working group
Internet Draft
Intended status: Standards Track
Expires: September 2011

Q. Zeng
J. Dong
Huawei Technologies
J. Heitz
Ericsson Inc.
K. Patel
Cisco Systems
R. Shakir
C&W

One-time Extended Community Based Outbound Route Filter for BGP-4

draft-dong-idr-one-time-ext-community-orf-00.txt

Abstract

This document defines a new Outbound Router Filter (ORF) type for BGP, termed "One-time Extended Community Outbound Route Filter", which would allow a BGP speaker to send to its BGP peer a route refresh request with a set of extended-community-based filters to make the peer re-advertise only the specific routes matching the filters to the speaker. This ORF-type enables a BGP speaker to refresh some specific routes without requiring its peer to re-advertise the whole Adj-RIB-Out, which makes the route refresh operation more efficient and reduces the impact on network stability. This filter does not change the outbound route filters on BGP peers and should only be used for one-time filtering.

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/ietf/lid-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 September 7, 2011.

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.

Conventions used in this document

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].

Table of Contents

1. Introduction	2
2. One-time Extended Community ORF-Type.....	3
3. Operation	4
4. Security Considerations.....	5
5. IANA Considerations	5
6. Acknowledgments	5
7. References	5
7.1. Normative References.....	5
7.2. Informative References.....	6
Authors' Addresses	7

1. Introduction

The Outbound Route Filtering Capability defined in [RFC5291] provides a mechanism for a BGP speaker to send to its BGP peer a set of Outbound Route Filters (ORFs) that can be used by its peer to filter its outbound routing updates to the speaker.

During some network operations, BGP speaker only needs to retrieve some routes with specific extended communities from its peer, but sending plain ROUTE-REFRESH will lead to the peer re-advertising its whole Adj-RIB-Out. Such large amounts of updates include a lot of unnecessary routes which would result in waste of processing resources and bandwidth. With the increase of IPv6 deployment, this problem could be more significant. Even configured with ORF mechanism as defined in [RFC5291], on receipt of a ROUTE-REFRESH message, the peer will re-advertise all the routes matching current outbound route filters, i.e., the whole Adj-Rib-Out for this BGP speaker. Since in this case the BGP speaker does not want to change the outbound route filters on its peer, this requirement cannot be met by current ORF mechanism.

This document defines a new Outbound Router Filter (ORF) type for BGP, termed "One-time Extended Community Outbound Route Filter", which would allow a BGP speaker to send to its BGP peer a route refresh request with a set of Extended Community based filters to make the peer re-advertise only the specific routes matching the filters to the speaker. This ORF-type enables a BGP speaker to retrieve routes with specific Extended Communities without requiring its peer to re-advertise the whole Adj-RIB-Out, which makes such route refresh operation more efficient and also reduces the impact on network stability. This filter does not change the outbound route filters on BGP peers and should only be used for one-time filtering.

One use case of one-time Extended Community ORF would be to refresh routes with specific Route Target (RT) Extended Community. For example, on receipt of routes with specific RTs, according to local policies some attributes of the routes may be changed, or some routes may be discarded. When later such local policies are changed or removed, the routes impacted by such policies need to be refreshed and processed according to the new local policies. With the whole Adj-RIB-Out route refresh it would result in a lot of unnecessary routes being re-advertised, and this would be a waste of the processing resource and bandwidth. In this case, one-time Extended Community ORF would be quite useful to request only routes matching specific RTs to be re-advertised.

2. One-time Extended Community ORF-Type

This document defines a new ORF type: One-time Extended Community ORF. Value of this ORF-Type is to be assigned by IANA.

In the following description, the sending speaker sends a one-time

ORF request and the receiving speaker receives it and sends back the routes to satisfy the request.

As specified in the [RFC5291], an ORF entry is a tuple of the form <AFI/SAFI, ORF-Type, Action, Match, ORF-value> an ORF consists of one or more ORF entries that have a common AFI/SAFI and ORF-Type. An ORF is identified by <AFI/SAFI, ORF-Type>.

The type-specific part consists of a single Extended Community encoded as an eight-octets field.

Since the semantics of this new ORF-Type is "one-time filtering" and has no impact on existing ORFs, the Action field is irrelevant and MUST be ignored on receipt.

The MATCH field of the One-time Extended Community ORF SHOULD be set to PERMIT on the sender and SHOULD be ignored on the receiver. This is the same as defined in Extended-Community ORF [EXT-COMM-ORF].

The ORF entries of this type would only be used as one-time filters that MUST not change any previously installed ORF entry on the receiving speaker.

3. Operation

The capability negotiation of <AFI/SAFI, One-time Extended Community ORF> MUST NOT delay the advertisement of routes with this AFI/SAFI.

The received One-time Extended Community ORF entries SHOULD only be used for one-time route filtering and MUST NOT be saved locally. The received One-time Extended Community ORF entries MUST NOT modify the outbound route filters on the receiving speaker (either locally configured or received from the sending speaker through ORF).

On receipt of ROUTE-REFRESH message with One-time Extended Community ORF entries, the receiving speaker SHOULD re-advertise to the sending speaker the routes from the Adj-RIB-Out associated with the sending speaker which pass the entries carried in the One-time Extended Community ORF as well as the locally saved ORFs (if any) received from the sending speaker.

Since different processing orders may lead to different results, the One-time ORFs and the regular ORFs SHOULD not be encoded in one route-refresh message.

During the period when the receiving speaker is sending updates to satisfy the One-time ORF request, it may experience other routing activity that will require it to send updates unrelated to the One-time ORF request. It is permitted to send these updates before it has completed sending the One-time ORF related updates.

Similarly, if a route that passes the One-time ORF has already been sent and the receiving speaker experiences routing activity that changes this route and the receiving speaker has not yet sent all routes to satisfy the One-time ORF request, it is permitted to send the changed route immediately.

Details about how to interoperate when both One-time ORF Capability and the Enhanced Route Refresh Capability as described in [Enhanced-Refresh] are enabled will be discussed in the next version.

4. Security Considerations

This extension to BGP does not change the underlying security issues in [RFC4271].

5. IANA Considerations

This document specifies a new Outbound Route Filtering (ORF) type, One-time Extended Community ORF. The value of the ORF-type needs to be assigned by IANA.

6. Acknowledgments

The authors would like to thank Robert Raszuk, John Scudder, Susan Hares, Haibo Wang, Jiawei Dong, Yaqun Xiao, Mach Chen for their valuable suggestions and comments to this document.

7. References

7.1. 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.
- [RFC5291] Chen, E. and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", RFC 5291, August 2008.

- [EXT-COMM-ORF] Chen, E., and Y. Rekhter, "Extended Community Based Outbound Route Filter for BGP-4", draft-chen-bgp-ext-community-orf-00, work in progress, June 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4020] Kompella, K. and A. Zinin, "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 4020, February 2005.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

7.2. Informative References

- [Enhanced-Refresh] K. Patel, E. Chen and B. Venkatachalapathy, "Enhanced Route Refresh Capability for BGP-4", draft-keyur-bgp-enhanced-route-refresh-01.txt, October 2010

Authors' Addresses

Jie Dong
Huawei Technologies Co.,Ltd.
Huawei Building, No.3 Xixi Rd.,
Hai-Dian District
Beijing, 100085
P.R. China

Email: jie.dong@huawei.com

Qing Zeng
Huawei Technologies Co.,Ltd.
Huawei Building, No.3 Xixi Rd.,
Hai-Dian District
Beijing, 100085
P.R. China

Email: zengqing@huawei.com

Jakob Heitz
Ericsson Inc.
100 Headquarters Drive
San Jose CA 95134
USA

Email: jakob.heitz@ericsson.com

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

Email: keyupate@cisco.com

Rob Shakir
Cable&Wireless Worldwide

Email: rob.shakir@cw.com

