

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2012

J. Dong
Q. Zeng
Huawei Technologies
J. Heitz
Ericsson Inc.
K. Patel
Cisco Systems
R. Shakir
BT
March 9, 2012

One-time Extended Community Based Outbound Route Filter for BGP-4
draft-dong-idr-one-time-ext-community-orf-02

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.

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](#)].

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

Internet-Draft

One-time Extended Community Based ORF

March 2012

material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 10, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/bcp78) 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.

Internet-Draft

One-time Extended Community Based ORF

March 2012

Table of Contents

1.	Introduction	4
2.	One-time Extended Community ORF-Type	5
3.	Operations	5
4.	Security Considerations	7
5.	IANA Considerations	7
6.	Acknowledgements	7
7.	Normative References	7
	Authors' Addresses	8

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, a 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 a large amount of updates includes 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 the ORF mechanism as defined in [[RFC5291](#)], on receipt of a ROUTE-REFRESH message, the peer will re-advertise all the routes matching the 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 the 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.

Another use case is network maintenance or verification. During maintenance, it may be revealed that some routes may be incorrect. A

refresh of a small segment of the routing table can help to correct those routes without requiring a refresh of the total routing table.

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 ORF entry consists of a single Extended Community, encoded as either 8 octets [[RFC4360](#)], or 20 octets [[RFC5701](#)]. The encoding of the type-specific part is as below:

```

+-----+
|   Ext-community Length (1 octet)   |
+-----+
|   Ext-community value (8 or 20 octets) |
+-----+

```

The "Ext-community Length" field contains length in octets of the extended community in the "Ext-community value" 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 [[I-D.chen-bgp-ext-community-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. Operations

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. A withdrawal of the route counts as a valid change.

If the receiving speaker has received the Enhanced Route Refresh Capability as described in [[I-D.ietf-idr-bgp-enhanced-route-refresh](#)], then it SHALL perform the following additional procedures.

It SHALL send a ROUTE-REFRESH message with the subtype to indicate demarcation of the beginning of route refresh (start-of-refresh) before sending routes to satisfy the ORF request and send a ROUTE-REFRESH message with the subtype to indicate demarcation of the ending of route refresh (end-of-refresh) after sending all routes to satisfy the request.

As part of the start-of-refresh and end-of-refresh messages, it SHALL include the one-time-ORF rules that it is satisfying with this refresh. The presence of the ORF is indicated in the same way as it is with the normal route refresh as in [[RFC5291](#)]: A BGP speaker can distinguish an incoming ROUTE-REFRESH message that carries one or

more ORF entries from an incoming plain ROUTE-REFRESH message by using the Message Length field in the BGP message header.

As with the procedures without the Enhanced Route Refresh Capability, the receiving speaker is permitted to send updates for routes unrelated to the ORF request before it has sent all updates to satisfy the current ORF request.

If a receiving speaker receives a new one-time-ORF request before it has finished completing a previous one, it MAY stop sending routes for the previous ORF request. It MUST NOT send an end-of-refresh for the previous request. If the speaker has received the Enhanced Route Refresh Capability, it SHALL send a start-of-refresh for the new request and start sending updates for it.

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. Acknowledgements

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. Normative References

[I-D.chen-bgp-ext-community-orf]
Chen, E., Lo, A., and K. Patel, "Extended Community Based Outbound Route Filter for BGP-4",
[draft-chen-bgp-ext-community-orf-02](#) (work in progress),
December 2011.

[I-D.ietf-idr-bgp-enhanced-route-refresh]
Patel, K., Chen, E., and B. Venkatachalapathy, "Enhanced Route Refresh Capability for BGP-4",
[draft-ietf-idr-bgp-enhanced-route-refresh-01](#) (work in

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2918] Chen, E., "Route Refresh Capability for BGP-4", [RFC 2918](#), September 2000.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), February 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC5291] Chen, E. and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", [RFC 5291](#), August 2008.
- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", [RFC 5701](#), November 2009.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Building, No.3 Xinxu Rd
Beijing 100085
China

Email: jie.dong@huawei.com

Qing Zeng
Huawei Technologies
Huawei Building, No.3 Xinxu Rd
Beijing 100085
China

Email: zengqing@huawei.com

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

Email: jakob.heizt@ericsson.com

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

Email: keyupate@cisco.com

Rob Shakir
BT
London
UK

Email: rob.shakir@bt.com

