

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2013

Q. Zeng

J. Dong
Huawei Technologies
J. Heitz
Ericsson Inc.
K. Patel
Cisco Systems
R. Shakir
BT
Z. Huang
China Telecom
October 22, 2012

One-time Address-Prefix Based Outbound Route Filter for BGP-4
draft-zeng-idr-one-time-prefix-orf-03

Abstract

This document defines a new Outbound Router Filter (ORF) type for BGP, termed "One-time Address Prefix Outbound Route Filter", which would allow a BGP speaker to send to its BGP peer a route refresh request with a set of address-prefix-based filters to make the peer send only the specific routes matching the filters to the speaker. This ORF-type enables a BGP speaker to re-advertise some specific routes without the need of advertising the whole Adj-RIB-Out of a specific address family, which makes the route recovery and trouble shooting operation more efficient and also reduces the impact on network stability. This filter does not change the outbound route filters or policies on the BGP peer and should only be used for one-time route 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 material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 25, 2013.

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](#) 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.

Table of Contents

1.	Introduction	4
2.	One-time Address Prefix ORF-Type	5
3.	Operation	5
4.	IANA Considerations	6
5.	Security Considerations	6
6.	Acknowledgements	6
7.	References	6
7.1.	Normative References	6
7.2.	Informative References	7
	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.

In some scenarios, BGP speaker only needs to retrieve some specific routes from its peer if the routes are possibly lost or contain some problematic attributes for some reason, but sending ROUTE-REFRESH message [[RFC2918](#)] will lead to the peer re-advertising its whole Adj-RIB-Out. Such large numbers of updates include a lot of unnecessary route updates which may make trouble shooting operation (such as packets tracking) more difficult, and is a waste of the processing resources and network 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 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 problem cannot be solved by existing ORF mechanism.

This document defines a new Outbound Router Filter (ORF) type for BGP, termed "One-time Address Prefix Outbound Route Filter", which would allow a BGP speaker to send to its BGP peer a route refresh request with a set of address-prefix-based filters to make the peer send only the specific routes matching these filters to the speaker. This new ORF-type enables a BGP speaker to re-advertise some specific routes without the need of advertising the whole Adj-RIB-Out of a particular address family, which makes the route recovery and trouble shooting operation (such as packet tracking) more efficient and also reduces the impact on network stability. This filter does not change the outbound route filters or policies on the BGP peer and should only be used for one-time route filtering.

Consider the following scenario: In an Inter-AS environment, ASBR-A received a malformed UPDATE from ASBR-B and treated it as withdraw according to [[I-D.ietf-idr-error-handling](#)]. While such event would be locally logged and the operators may be notified, it is important for ASBR-A to try to recover these routes as soon as possible since the routes which are treated as withdraw may impact some critical services. A good method is to ask the peering ASBR-B to re-advertise such routes with some back off mechanism. One-time Prefix ORF is a low impact way to achieve this.

2. One-time Address Prefix ORF-Type

This document defines a new ORF type: One-time Address Prefix ORF.

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 format of One-time Address Prefix ORF-Type entry is the same as the encoding of Address Prefix ORF in [[RFC5292](#)], with the specific fields defined as follows:

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

The matching rules of the One-time Address Prefix ORF are the same as defined in Address-Prefix-Based ORF [[RFC5292](#)].

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

3. Operation

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

The received One-time Address Prefix ORF entries SHOULD only be used for one-time route filtering and MUST NOT be saved locally. The received One-time Address Prefix 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 Address Prefix 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 Address Prefix 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 [[I-D.ietf-idr-bgp-enhanced-route-refresh](#)] are enabled will be discussed in a future version.

4. IANA Considerations

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

5. Security Considerations

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

6. Acknowledgements

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

7. References

7.1. Normative References

[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.
- [RFC5291] Chen, E. and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", [RFC 5291](#), August 2008.
- [RFC5292] Chen, E. and S. Sangli, "Address-Prefix-Based Outbound Route Filter for BGP-4", [RFC 5292](#), August 2008.

7.2. Informative References

- [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-02](#) (work in progress), June 2012.
- [I-D.ietf-idr-error-handling]
Scudder, J., Chen, E., Mohapatra, P., and K. Patel,
"Revised Error Handling for BGP UPDATE Messages",
[draft-ietf-idr-error-handling-02](#) (work in progress),
June 2012.

Authors' Addresses

Qing Zeng
Beijing
China

Email: zengqqqq@gmail.com

Jie Dong
Huawei Technologies
Huawei Building, No.156 Beiqing Rd
Beijing 100095
China

Email: jie.dong@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

ZhiLan Huang
China Telecom
109 West Zhongshan Ave
Guangzhou 510630
China

Email: huangzl@gsta.com

