

Internet Engineering Task Force (IETF)
Request for Comments: 7947
Category: Standards Track
ISSN: 2070-1721

E. Jasinska
BigWave IT
N. Hilliard
INEX
R. Raszuk
Bloomberg LP
N. Bakker
Akamai Technologies B.V.
September 2016

Internet Exchange BGP Route Server

Abstract

This document outlines a specification for multilateral interconnections at Internet Exchange Points (IXPs). Multilateral interconnection is a method of exchanging routing information among three or more External BGP (EBGP) speakers using a single intermediate broker system, referred to as a route server. Route servers are typically used on shared access media networks, such as IXPs, to facilitate simplified interconnection among multiple Internet routers.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7947>.

Copyright Notice

Copyright (c) 2016 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 to Multilateral Interconnection	3
1.1.	Notational Conventions	3
2.	Technical Considerations for Route Server Implementations . .	4
2.1.	Client UPDATE Messages	4
2.2.	Attribute Transparency	4
2.2.1.	NEXT_HOP Attribute	4
2.2.2.	AS_PATH Attribute	5
2.2.2.1.	Route Server AS_PATH Management	5
2.2.2.2.	Route Server client AS_PATH Management	5
2.2.3.	MULTI_EXIT_DISC Attribute	5
2.2.4.	Communities Attributes	5
2.3.	Per-Client Policy Control in Multilateral Interconnection	6
2.3.1.	Path Hiding on a Route Server	6
2.3.2.	Mitigation of Path Hiding	7
2.3.2.1.	Multiple Route Server RIBs	7
2.3.2.2.	Advertising Multiple Paths	8
2.3.3.	Implementation Suggestions	9
3.	Security Considerations	9
4.	References	10
4.1.	Normative References	10
4.2.	Informative References	10
	Acknowledgments	11
	Authors' Addresses	12

1. Introduction to Multilateral Interconnection

Internet Exchange Points (IXPs) provide IP data interconnection facilities for their participants, typically using shared Layer 2 networking media such as Ethernet. The Border Gateway Protocol (BGP) [RFC4271], an inter-Autonomous System (inter-AS) routing protocol, is commonly used to facilitate exchange of network reachability information over such media.

While bilateral EBGp sessions between exchange participants were previously the most common means of exchanging reachability information, the overhead associated with dense interconnection can cause substantial operational scaling problems for participants of larger IXPs.

Multilateral interconnection is a method of interconnecting BGP speaking routers using a third-party brokering system, commonly referred to as a route server and typically managed by the IXP operator. Each multilateral interconnection participant (usually referred to as a "route server client") announces network reachability information to the route server using EBGp. The route server, in turn, forwards this information to each route server client connected to it, according to its configuration. Although a route server uses BGP to exchange reachability information with each of its clients, it does not forward traffic itself and is therefore not a router.

A route server can be viewed as similar in function to a route reflector [RFC4456], except that it operates using EBGp instead of Internal BGP (IBGP). Certain adaptations to [RFC4271] are required to enable an EBGp router to operate as a route server; these are outlined in Section 2 of this document. Route server functionality is not mandatory in BGP implementations.

The term "route server" is often used in a different context to describe a BGP node whose purpose is to accept BGP feeds from multiple clients for the purpose of operational analysis and troubleshooting. A system of this form may alternatively be known as a "route collector" or a "route-views server". This document uses the term "route server" exclusively to describe multilateral peering brokerage systems.

1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Technical Considerations for Route Server Implementations

A route server uses BGP [RFC4271] to broker network reachability information amongst its clients. There are some differences between the behavior of a BGP route server and a BGP implementation that is strictly compliant with [RFC4271]. These differences are described as follows.

2.1. Client UPDATE Messages

A route server MUST accept all UPDATE messages received from each of its clients for inclusion in its Adj-RIB-In. These UPDATE messages MAY be omitted from the route server's Loc-RIB or Loc-RIBs, due to filters configured for the purpose of implementing routing policy. The route server SHOULD perform one or more BGP Decision Processes to select routes for subsequent advertisement to its clients, taking into account possible configuration to provide multiple Network Layer Reachability Information (NLRI) paths to a particular client as described in Section 2.3.2.2 or multiple Loc-RIBs as described in Section 2.3.2.1. The route server SHOULD forward UPDATE messages from its Loc-RIB or Loc-RIBs to its clients as determined by local policy.

2.2. Attribute Transparency

As a route server primarily performs a brokering service, modification of attributes could cause route server clients to alter their BGP Decision Process for received prefix reachability information, thereby changing the intended routing policies of exchange participants. Therefore, contrary to what is specified in Section 5 of [RFC4271], route servers SHOULD NOT by default (unless explicitly configured) update well-known BGP attributes received from route server clients before redistributing them to their other route server clients. Optional recognized and unrecognized BGP attributes, whether transitive or non-transitive, SHOULD NOT be updated by the route server (unless enforced by local IXP operator configuration) and SHOULD be passed on to other route server clients.

2.2.1. NEXT_HOP Attribute

The NEXT_HOP is a well-known mandatory BGP attribute that defines the IP address of the router used as the next hop to the destinations listed in the NLRI field of the UPDATE message. As the route server does not participate in the actual routing of traffic, the NEXT_HOP attribute MUST be passed unmodified to the route server clients, similar to the "third-party" next-hop feature described in Section 5.1.3. of [RFC4271].

2.2.2. AS_PATH Attribute

AS_PATH is a well-known mandatory attribute that identifies the ASes through which routing information carried in the UPDATE message has passed.

2.2.2.1. Route Server AS_PATH Management

As a route server does not participate in the process of forwarding data between client routers, and because modification of the AS_PATH attribute could affect the route server client BGP Decision Process, the route server SHOULD NOT prepend its own AS number to the AS_PATH segment nor modify the AS_PATH segment in any other way. This differs from the behavior specified in Section 5.1.2 of [RFC4271], which requires that the BGP speaker prepends its own AS number as the last element of the AS_PATH segment. This is a recommendation rather than a requirement solely to provide backwards compatibility with legacy route server client implementations that do not yet support the requirements specified in Section 2.2.2.2.

2.2.2.2. Route Server client AS_PATH Management

In contrast to what is recommended in Section 6.3 of [RFC4271], route server clients need to be able to accept UPDATE messages where the leftmost AS in the AS_PATH attribute is not equal to the AS number of the route server that sent the UPDATE message. If the route server client BGP system has implemented a check for this, the BGP implementation MUST allow this check to be disabled and SHOULD allow the check to be disabled on a per-peer basis.

2.2.3. MULTI_EXIT_DISC Attribute

MULTI_EXIT_DISC is an optional non-transitive attribute intended to be used on external (inter-AS) links to discriminate among multiple exit or entry points to the same neighboring AS. Contrary to Section 5.1.4 of [RFC4271], if applied to an NLRI UPDATE sent to a route server, this attribute SHOULD be propagated to other route server clients, and the route server SHOULD NOT modify its value.

2.2.4. Communities Attributes

The BGP Communities [RFC1997] and Extended Communities [RFC4360] attributes are intended for labeling information carried in BGP UPDATE messages. Transitive as well as non-transitive Communities attributes applied to an NLRI UPDATE sent to a route server SHOULD NOT be modified, processed, or removed, except as defined by local

policy. If a Communities attribute is intended for processing by the route server itself, as determined by local policy, it MAY be modified or removed.

2.3. Per-Client Policy Control in Multilateral Interconnection

While IXP participants often use route servers with the intention of interconnecting with as many other route server participants as possible, there are circumstances where control of path distribution on a per-client basis is important to ensure that desired interconnection policies are met.

The control of path distribution on a per-client basis can lead to a path being hidden from the route server client. We refer to this as "path hiding".

Neither Section 2.3 nor its subsections form part of the normative specification of this document; they are included for information purposes only.

2.3.1. Path Hiding on a Route Server




Figure 1: Per-Client Policy Controlled Interconnection at an IXP

Using the example in Figure 1, AS1 does not directly exchange prefix information with either AS2 or AS3 at the IXP but only interconnects with AS4. The lines between AS1, AS2, AS3, and AS4 represent interconnection relationships, whether via bilateral or multilateral connections.

In the traditional bilateral interconnection model, per-client policy control to a third-party exchange participant is accomplished either by not engaging in a bilateral interconnection with that participant or by implementing outbound filtering on the BGP session towards that

participant. However, in a multilateral interconnection environment, only the route server can perform outbound filtering in the direction of the route server client; route server clients depend on the route server to perform their outbound filtering for them.

Assuming the BGP Decision Process [RFC4271] is used, when the same prefix is advertised to a route server from multiple route server clients, the route server will select a single path for propagation to all connected clients. If, however, the route server has been configured to filter the calculated best path from reaching a particular route server client, then that client will not receive a path for that prefix, although alternate paths received by the route server might have been policy compliant for that client. This phenomenon is referred to as "path hiding".

For example, in Figure 1, if the same prefix were sent to the route server via AS2 and AS4, and the route via AS2 was preferred according to the BGP Decision Process on the route server, but AS2's policy prevented the route server from sending the path to AS1, then AS1 would never receive a path to this prefix, even though the route server had previously received a valid alternative path via AS4. This happens because the BGP Decision Process is performed only once on the route server for all clients.

Path hiding will only occur on route servers that employ per-client policy control; if an IXP operator deploys a route server without implementing a per-client routing policy control system, then path hiding does not occur, as all paths are considered equally valid from the point of view of the route server.

2.3.2. Mitigation of Path Hiding

There are several approaches that can be taken to mitigate against path hiding.

2.3.2.1. Multiple Route Server RIBs

The most portable method to allow for per-client policy control without the occurrence of path hiding is to use a route server BGP implementation that performs the per-client best path calculation for each set of paths to a prefix, which results after the route server's client policies have been taken into consideration. This can be implemented by using per-client Loc-RIBs, with path filtering implemented between the Adj-RIB-In and the per-client Loc-RIB. Implementations can optimize this by maintaining paths not subject to filtering policies in a global Loc-RIB, with per-client Loc-RIBs stored as deltas.

This implementation is highly portable, as it makes no assumptions about the feature capabilities of the route server clients.

2.3.2.2. Advertising Multiple Paths

The path distribution model described above assumes standard BGP session encoding where the route server sends a single path to its client for any given prefix. This path is selected using the BGP path selection Decision Process described in [RFC4271]. If, however, it were possible for the route server to send more than a single path to a route server client, then route server clients would no longer depend on receiving a single path to a particular prefix; consequently, the path-hiding problem described in Section 2.3.1 would disappear.

We present two methods that describe how such increased path diversity could be implemented.

2.3.2.2.1. Diverse BGP Path Approach

The diverse BGP path proposal as defined in [RFC6774] is a simple way to distribute multiple prefix paths from a route server to a route server client by using a separate BGP session from the route server to a client for each different path.

The number of paths that may be distributed to a client is constrained by the number of BGP sessions that the server and the client are willing to establish with each other. The distributed paths may be established from the global BGP Loc-RIB on the route server in addition to any per-client Loc-RIB. As there may be more potential paths to a given prefix than configured BGP sessions, this method is not guaranteed to eliminate the path-hiding problem in all situations. Furthermore, this method may significantly increase the number of BGP sessions handled by the route server, which may negatively impact its performance.

2.3.2.2.2. BGP ADD-PATH Approach

[RFC7911] proposes a different approach to multiple path propagation, by allowing a BGP speaker to forward multiple paths for the same prefix on a single BGP session. As [RFC4271] specifies that a BGP listener must implement an implicit withdraw when it receives an UPDATE message for a prefix that already exists in its Adj-RIB-In, this approach requires explicit support for the feature both on the route server and on its clients.

If the ADD-PATH capability is negotiated bidirectionally between the route server and a route server client, and the route server client propagates multiple paths for the same prefix to the route server, then this could potentially cause the propagation of inactive, invalid, or suboptimal paths to the route server, thereby causing loss of reachability to other route server clients. For this reason, ADD-PATH implementations on a route server should enforce a send-only mode with the route server clients, which would result in negotiating a receive-only mode from the client to the route server.

2.3.3. Implementation Suggestions

Authors of route server implementations may wish to consider one of the methods described in Section 2.3.2 to allow per-client route server policy control without path hiding.

Recommendations for route server operations are described separately in [RFC7948].

3. Security Considerations

The path-hiding problem outlined in Section 2.3.1 can be used in certain circumstances to proactively block third-party path announcements from other route server clients. Route server operators should be aware that security issues may arise unless steps are taken to mitigate against path hiding.

The AS_PATH check described in Section 2.2.2 is normally enabled in order to check for malformed AS paths. If this check is disabled, the route server client loses the ability to check incoming UPDATE messages for certain categories of problems. This could potentially cause corrupted BGP UPDATE messages to be propagated where they might not be propagated if the check were enabled. Regardless of any problems relating to malformed UPDATE messages, this check is also used to detect BGP loops; removing the check could potentially cause routing loops to be formed. Consequently, this check SHOULD NOT be disabled by IXP participants unless it is needed to establish BGP sessions with a route server and, if possible, should only be disabled for peers that are route servers.

Route server operators should carefully consider the security practices discussed in "BGP Operations and Security" [RFC7454].

4. References

4.1. Normative References

- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<http://www.rfc-editor.org/info/rfc1997>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<http://www.rfc-editor.org/info/rfc4271>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<http://www.rfc-editor.org/info/rfc4360>>.

4.2. Informative References

- [RFC1863] Haskin, D., "A BGP/IDRP Route Server alternative to a full mesh routing", RFC 1863, DOI 10.17487/RFC1863, October 1995, <<http://www.rfc-editor.org/info/rfc1863>>.
- [RFC4223] Savola, P., "Reclassification of RFC 1863 to Historic", RFC 4223, DOI 10.17487/RFC4223, October 2005, <<http://www.rfc-editor.org/info/rfc4223>>.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<http://www.rfc-editor.org/info/rfc4456>>.
- [RFC6774] Raszuk, R., Ed., Fernando, R., Patel, K., McPherson, D., and K. Kumaki, "Distribution of Diverse BGP Paths", RFC 6774, DOI 10.17487/RFC6774, November 2012, <<http://www.rfc-editor.org/info/rfc6774>>.
- [RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <<http://www.rfc-editor.org/info/rfc7454>>.

- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<http://www.rfc-editor.org/info/rfc7911>>.
- [RFC7948] Hilliard, N., Jasinska, E., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server Operations", RFC 7948, DOI 10.17487/RFC7948, September 2016, <<http://www.rfc-editor.org/info/rfc7948>>.

Acknowledgments

The authors would like to thank Ryan Bickhart, Steven Bakker, Martin Pels, Chris Hall, Aleksu Suhonen, Bruno Decraene, Pierre Francois, and Eduardo Ascenco Reis for their valuable input.

In addition, the authors would like to acknowledge the developers of BIRD, OpenBGPD, Quagga, and IOS whose BGP implementations include route server capabilities that are compliant with this document.

Route server functionality was described in 1995 in [RFC1863], and modern route server implementations are based on concepts developed in the 1990s by the Routing Arbiter Project and the Route Server Next Generation (RSNG) Project, managed by ISI and Merit. Although the original RSNG code is no longer in use at any IXPs, the IXP community owes a debt of gratitude to the many people who were involved in route server development in the 1990s. Note that [RFC1863] was made historical by [RFC4223].

Authors' Addresses

Elisa Jasinska
BigWave IT
ul. Skawinska 27/7
Krakow, MP 31-066
Poland

Email: elisa@bigwaveit.org

Nick Hilliard
INEX
4027 Kingswood Road
Dublin 24
Ireland

Email: nick@inex.ie

Robert Raszuk
Bloomberg LP
731 Lexington Ave
New York City, NY 10022
United States of America

Email: robert@raszuk.net

Niels Bakker
Akamai Technologies B.V.
Kingsfordweg 151
Amsterdam 1043 GR
Netherlands

Email: nbakker@akamai.com