

Internet Engineering Task Force (IETF)
Request for Comments: 7948
Category: Informational
ISSN: 2070-1721

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

Internet Exchange BGP Route Server Operations

Abstract

The popularity of Internet Exchange Points (IXPs) brings new challenges to interconnecting networks. While bilateral External BGP (EBGP) sessions between exchange participants were historically the most common means of exchanging reachability information over an IXP, the overhead associated with this interconnection method causes serious operational and administrative scaling problems for IXP participants.

Multilateral interconnection using Internet route servers can dramatically reduce the administrative and operational overhead associated with connecting to IXPs; in some cases, route servers are used by IXP participants as their preferred means of exchanging routing information.

This document describes operational considerations for multilateral interconnections at IXPs.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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/rfc7948>.

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	3
1.1. Notational Conventions	3
2. Bilateral BGP Sessions	3
3. Multilateral Interconnection	4
4. Operational Considerations for Route Server Installations . .	6
4.1. Path Hiding	6
4.2. Route Server Scaling	6
4.2.1. Tackling Scaling Issues	7
4.2.1.1. View Merging and Decomposition	7
4.2.1.2. Destination Splitting	8
4.2.1.3. NEXT_HOP Resolution	8
4.3. Prefix Leakage Mitigation	8
4.4. Route Server Redundancy	9
4.5. AS_PATH Consistency Check	9
4.6. Export Routing Policies	10
4.6.1. BGP Communities	10
4.6.2. Internet Routing Registries	10
4.6.3. Client-Accessible Databases	10
4.7. Layer 2 Reachability Problems	11
4.8. BGP NEXT_HOP Hijacking	11
4.9. BGP Operations and Security	13
5. Security Considerations	13
6. References	13
6.1. Normative References	13
6.2. Informative References	14
Acknowledgments	15
Authors' Addresses	15

1. Introduction

Internet Exchange Points (IXPs) provide IP data interconnection facilities for their participants, using data link-layer protocols such as Ethernet. The Border Gateway Protocol (BGP) [RFC4271] is normally used to facilitate exchange of network reachability information over these media.

As bilateral interconnection between IXP participants requires operational and administrative overhead, BGP route servers [RFC7947] are often deployed by IXP operators to provide a simple and convenient means of interconnecting IXP participants with each other. A route server redistributes BGP routes received from its BGP clients to other clients according to a prespecified policy, and it can be viewed as similar to an EBGPP equivalent of an Internal BGP (IBGP) [RFC4456] route reflector.

Route servers at IXPs require careful management, and it is important for route server operators to thoroughly understand both how they work and what their limitations are. In this document, we discuss several issues of operational relevance to route server operators and provide recommendations to help route server operators provision a reliable interconnection service.

1.1. Notational Conventions

The keywords "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].

The phrase "BGP route" in this document should be interpreted as the term "Route" described in [RFC4271].

2. Bilateral BGP Sessions

Bilateral interconnection is a method of interconnecting routers using individual BGP sessions between each pair of participant routers on an IXP, in order to exchange reachability information. If an IXP participant wishes to implement an open interconnection policy -- i.e., a policy of interconnecting with as many other IXP participants as possible -- it is necessary for the participant to liaise with each of their intended interconnection partners. Interconnection can then be implemented bilaterally by configuring a BGP session on both participants' routers to exchange network reachability information. If each exchange participant interconnects with each other participant, a full mesh of BGP sessions is needed, as shown in Figure 1.




Figure 1: Full-Mesh Interconnection at an IXP

Figure 1 depicts an IXP platform with four connected routers, administered by four separate exchange participants, each of them with a locally unique Autonomous System (AS) number: AS1, AS2, AS3, and AS4. The lines between the routers depict BGP sessions; the dotted edge represents the IXP border. Each of these four participants wishes to exchange traffic with all other participants; this is accomplished by configuring a full mesh of BGP sessions on each router connected to the exchange, resulting in six BGP sessions across the IXP fabric.

The number of BGP sessions at an exchange has an upper bound of $n*(n-1)/2$, where n is the number of routers at the exchange. As many exchanges have large numbers of participating networks, the amount of administrative and operation overhead required to implement an open interconnection scales quadratically. New participants to an IXP require significant initial resourcing in order to gain value from their IXP connection, while existing exchange participants need to commit ongoing resources in order to benefit from interconnecting with these new participants.

3. Multilateral Interconnection

Multilateral interconnection is implemented using a route server configured to distribute BGP routes among client routers. The route server preserves the BGP NEXT_HOP attribute from all received BGP routes and passes them with unchanged NEXT_HOP to its route server clients according to its configured routing policy, as described in [RFC7947]. Using this method of exchanging BGP routes, an IXP participant router can receive an aggregated list of BGP routes from all other route server clients using a single BGP session to the route server instead of depending on BGP sessions with each router at the exchange. This reduces the overall number of BGP sessions at an

Internet exchange from $n*(n-1)/2$ to n , where n is the number of routers at the exchange.

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.

In practical terms, this allows dense interconnection between IXP participants with low administrative overhead and significantly simpler and smaller router configurations. In particular, new IXP participants benefit from immediate and extensive interconnection, while existing route server participants receive reachability information from these new participants without necessarily having to modify their configurations.




Figure 2: IXP-Based Interconnection with Route Server

As illustrated in Figure 2, each router on the IXP fabric requires only a single BGP session to the route server, from which it can receive reachability information for all other routers on the IXP that also connect to the route server.

Multilateral and bilateral interconnections between different autonomous systems are not exclusive to each other, and it is not unusual to have both sorts of sessions configured in parallel at an IXP. This configuration will lead to additional paths being available to the BGP Decision Process, which will calculate a best path as normal.

4. Operational Considerations for Route Server Installations

4.1. Path Hiding

"Path hiding" is a term used in [RFC7947] to describe the process whereby a route server may mask individual paths by applying conflicting routing policies to its Loc-RIB. When this happens, route server clients receive incomplete information from the route server about network reachability.

There are several approaches that may be used to mitigate against the effect of path hiding; these are described in [RFC7947]. However, the only method that does not require explicit support from the route server client is for the route server itself to maintain an individual Loc-RIB for each client that is the subject of conflicting routing policies.

4.2. Route Server Scaling

While deployment of multiple Loc-RIBs on the route server presents a simple way to avoid the path-hiding problem noted in Section 4.1, this approach requires significantly more computing resources on the route server than where a single Loc-RIB is deployed for all clients. As the BGP Decision Process [RFC4271] must be applied to all Loc-RIBs deployed on the route server, both CPU and memory requirements on the host computer scale approximately according to $O(P * N)$, where P is the total number of unique paths received by the route server, and N is the number of route server clients that require a unique Loc-RIB. As this is a super-linear scaling relationship, large route servers may derive benefit from deploying per-client Loc-RIBs only where they are required.

Regardless of whether any Loc-RIB optimization technique is implemented, the route server's theoretical upper-bound network bandwidth requirements will scale according to $O(P_{tot} * N)$, where P_{tot} is the total number of unique paths received by the route server, and N is the total number of route server clients. In the case where P_{avg} (the arithmetic mean number of unique paths received per route server client) remains roughly constant even as the number of connected clients increases, the total number of prefixes will equal the average number of prefixes multiplied by the number of clients. Symbolically, this can be written as $P_{tot} = P_{avg} * N$. If we assume that in the worst case, each prefix is associated with a different set of BGP path attributes, so must be transmitted individually, the network bandwidth scaling function can be rewritten as $O((P_{avg} * N) * N)$ or $O(N^2)$. This quadratic upper bound on the network traffic requirements indicates that the route server model may not scale well for larger numbers of clients.

In practice, most prefixes will be associated with a limited number of BGP path attribute sets, allowing more efficient transmission of BGP routes from the route server than the theoretical analysis suggests. In the analysis above, P_{tot} will increase monotonically according to the number of clients, but it will have an upper limit of the size of the full default-free routing table of the network in which the IXP is located. Observations from production route servers have shown that most route server clients generally avoid using custom routing policies, and consequently, the route server may not need to deploy per-client Loc-RIBs. These practical bounds reduce the theoretical worst-case scaling scenario to the point where route server deployments are manageable even on larger IXPs.

4.2.1. Tackling Scaling Issues

The problem of scaling route servers still presents serious practical challenges and requires careful attention. Scaling analysis indicates problems in three key areas: route processor CPU overhead associated with BGP Decision Process calculations, the memory requirements for handling many different BGP path entries, and the network traffic bandwidth required to distribute these BGP routes from the route server to each route server client.

4.2.1.1. View Merging and Decomposition

View merging and decomposition, outlined in [RS-ARCH], describes a method of optimizing memory and CPU requirements where multiple route server clients are subject to exactly the same routing policies. In this situation, multiple Loc-RIB views can be merged into a single view.

There are several variations of this approach. If the route server operator has prior knowledge of interconnection relationships between route server clients, then the operator may configure separate Loc-RIBs only for route server clients with unique routing policies. As this approach requires prior knowledge of interconnection relationships, the route server operator must depend on each client sharing their interconnection policies either in an internal provisioning database controlled by the operator or in an external data store such as an Internet Routing Registry Database.

Conversely, the route server implementation itself may implement internal view decomposition by creating virtual Loc-RIBs based on a single in-memory master Loc-RIB, with delta differences for each prefix subject to different routing policies. This allows a more fine-grained and flexible approach to the problem of Loc-RIB scaling, at the expense of requiring a more complex in-memory Loc-RIB structure.

Whatever method of view merging and decomposition is chosen on a route server, pathological edge cases can be created whereby they will scale no better than fully non-optimized per-client Loc-RIBs. However, as most route server clients connect to a route server for the purposes of reducing overhead, rather than implementing complex per-client routing policies, edge cases tend not to arise in practice.

4.2.1.2. Destination Splitting

Destination splitting, also described in [RS-ARCH], describes a method for route server clients to connect to multiple route servers and to send non-overlapping sets of prefixes to each route server. As each route server computes the best path for its own set of prefixes, the quadratic scaling requirement operates on multiple smaller sets of prefixes. This reduces the overall computational and memory requirements for managing multiple Loc-RIBs and performing the best-path calculation on each.

In practice, the route server operator would need all route server clients to send a full set of BGP routes to each route server. The route server operator could then selectively filter these prefixes for each route server by using either BGP Outbound Route Filtering [RFC5291] or inbound prefix filters configured on client BGP sessions.

4.2.1.3. NEXT_HOP Resolution

As route servers are usually deployed at IXPs where all connected routers are on the same Layer 2 broadcast domain, recursive resolution of the NEXT_HOP attribute is generally not required and can be replaced by a simple check to ensure that the NEXT_HOP value for each received BGP route is a network address on the IXP LAN's IP address range.

4.3. Prefix Leakage Mitigation

Prefix leakage occurs when a BGP client unintentionally distributes BGP routes to one or more neighboring BGP routers. Prefix leakage of this form to a route server can cause serious connectivity problems at an IXP if each route server client is configured to accept all BGP routes from the route server. It is therefore RECOMMENDED when deploying route servers that, due to the potential for collateral damage caused by BGP route leakage, route server operators deploy prefix leakage mitigation measures in order to prevent unintentional prefix announcements or else limit the scale of any such leak. Although not foolproof, per-client inbound prefix limits can restrict the damage caused by prefix leakage in many cases. Per-client

inbound prefix filtering on the route server is a more deterministic and usually more reliable means of preventing prefix leakage but requires more administrative resources to maintain properly.

If a route server operator implements per-client inbound prefix filtering, then it is RECOMMENDED that the operator also builds in mechanisms to automatically compare the Adj-RIB-In received from each client with the inbound prefix lists configured for those clients. Naturally, it is the responsibility of the route server client to ensure that their stated prefix list is compatible with what they announce to an IXP route server. However, many network operators do not carefully manage their published routing policies, and it is not uncommon to see significant variation between the two sets of prefixes. Route server operator visibility into this discrepancy can provide significant advantages to both operator and client.

4.4. Route Server Redundancy

As the purpose of an IXP route server implementation is to provide a reliable reachability brokerage service, it is RECOMMENDED that exchange operators who implement route server systems provision multiple route servers on each shared Layer 2 domain. There is no requirement to use the same BGP implementation or operating system for each route server on the IXP fabric; however, it is RECOMMENDED that where an operator provisions more than a single server on the same shared Layer 2 domain, each route server implementation be configured equivalently and in such a manner that the path reachability information from each system is identical.

4.5. AS_PATH Consistency Check

[RFC4271] requires that every BGP speaker that advertises a BGP route to another external BGP speaker prepends its own AS number as the last element of the AS_PATH sequence. Therefore, the leftmost AS in an AS_PATH attribute should be equal to the AS number of the BGP speaker that sent the BGP route.

As [RFC7947] suggests that route servers should not modify the AS_PATH attribute, a consistency check on the AS_PATH of a BGP route received by a route server client would normally fail. It is therefore RECOMMENDED that route server clients disable the AS_PATH consistency check towards the route server.

4.6. Export Routing Policies

Policy filtering is commonly implemented on route servers to provide prefix distribution control mechanisms for route server clients. A route server "export" policy is a policy that affects prefixes sent from the route server to a route server client. Several different strategies are commonly used for implementing route server export policies.

4.6.1. BGP Communities

Prefixes sent to the route server are tagged with specific standard BGP Communities [RFC1997] or Extended Communities [RFC4360] attributes, based on predefined values agreed between the operator and all clients. Based on these Communities values, BGP routes may be propagated to all other clients, a subset of clients, or none. This mechanism allows route server clients to instruct the route server to implement per-client export routing policies.

As both standard BGP Communities and Extended Communities values are restricted to 6 octets or fewer, it is not possible for both the global and local administrator fields in the BGP Communities value to fit a 4-octet AS number. Bearing this in mind, the route server operator SHOULD take care to ensure that the predefined BGP Communities values mechanism used on their route server is compatible with 4-octet AS numbers [RFC6793].

4.6.2. Internet Routing Registries

Internet Routing Registry databases (IRRDBs) may be used by route server operators to construct per-client routing policies. "Routing Policy Specification Language (RPSL)" [RFC2622] provides a comprehensive grammar for describing interconnection relationships, and several toolsets exist that can be used to translate RPSL policy description into route server configurations.

4.6.3. Client-Accessible Databases

Should the route server operator not wish to use either BGP Communities or the public IRRDBs for implementing client export policies, they may implement their own routing policy database system for managing their clients' requirements. A database of this form SHOULD allow a route server client operator to update their routing policy and provide a mechanism for allowing the client to specify whether they wish to exchange all their prefixes with any other route server client. Optionally, the implementation may allow a client to specify unique routing policies for individual prefixes over which they have routing policy control.

4.7. Layer 2 Reachability Problems

Layer 2 reachability problems on an IXP can cause serious operational problems for IXP participants that depend on route servers for interconnection. Ethernet switch forwarding bugs have occasionally been observed to cause non-transitive reachability. For example, given a route server and two IXP participants, A and B, if the two participants can reach the route server but cannot reach each other, then traffic between the participants may be dropped until such time as the Layer 2 forwarding problem is resolved. This situation does not tend to occur in bilateral interconnection arrangements, as the routing control path between the two hosts is usually (but not always, due to IXP inter-switch connectivity load-balancing algorithms) the same as the data path between them.

Problems of this form can be partially mitigated by using Bidirectional Forwarding Detection (BFD) [RFC5881]. However, as this is a bilateral protocol configured between routers, and as there is currently no protocol to automatically configure BFD sessions between route server clients, BFD does not currently provide an optimal means of handling the problem. Even if automatic BFD session configuration were possible, practical problems would remain. If two IXP route server clients were configured to run BFD between each other and the protocol detected a non-transitive loss of reachability between them, each of those routers would internally mark the other's prefixes as unreachable via the BGP path announced by the route server. As the route server only propagates a single best path to each client, this could cause either sub-optimal routing or complete connectivity loss if there were no alternative paths learned from other BGP sessions.

4.8. BGP NEXT_HOP Hijacking

Item 2 in Section 5.1.3 of [RFC4271] allows EBGp speakers to change the NEXT_HOP address of a received BGP route to be a different Internet address on the same subnet. This is the mechanism that allows route servers to operate on a shared Layer 2 IXP network. However, the mechanism can be abused by route server clients to redirect traffic for their prefixes to other IXP participant routers.




Figure 3: BGP NEXT_HOP Hijacking Using a Route Server

For example, in Figure 3, if AS1 and AS2 both announce BGP routes for AS99 to the route server, AS1 could set the NEXT_HOP address for AS99's routes to be the address of AS2's router, thereby diverting traffic for AS99 via AS2. This may override the routing policies of AS99 and AS2.

Worse still, if the route server operator does not use inbound prefix filtering, AS1 could announce any arbitrary prefix to the route server with a NEXT_HOP address of any other IXP participant. This could be used as a denial-of-service mechanism against either the users of the address space being announced by illicitly diverting their traffic or the other IXP participant by overloading their network with traffic that would not normally be sent there.

This problem is not specific to route servers, and it can also be implemented using bilateral BGP sessions. However, the potential damage is amplified by route servers because a single BGP session can be used to affect many networks simultaneously.

Because route server clients cannot easily implement next-hop policy checks against route server BGP sessions, route server operators SHOULD check that the BGP NEXT_HOP attribute for BGP routes received from a route server client matches the interface address of the client. If the route server receives a BGP route where these addresses are different and where the announcing route server client is in a different AS to the route server client that uses the next-hop address, the BGP route SHOULD be dropped. Permitting next-hop

rewriting for the same AS allows an organization with multiple connections into an IXP configured with different IP addresses to direct traffic off the IXP infrastructure through any of their connections for traffic engineering or other purposes.

4.9. BGP Operations and Security

BGP route servers SHOULD be configured and operated in compliance with [RFC7454] with the exception of Section 11, "BGP Community Scrubbing", which may not necessarily apply on a route server, depending on the route server operator policy.

5. Security Considerations

On route server installations that do not employ path-hiding mitigation techniques, the path-hiding problem outlined in Section 4.1 could be used by an IXP participant to prevent the route server from sending any BGP routes for a particular prefix to other route server clients, even if there was a valid path to that destination via another route server client.

If the route server operator does not implement prefix leakage mitigation as described in Section 4.3, it is trivial for route server clients to implement denial-of-service attacks against arbitrary Internet networks by leaking BGP routes to a route server.

Route server installations SHOULD be secured against BGP NEXT_HOP hijacking, as described in Section 4.8.

6. References

6.1. Normative References

- [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>>.
- [RFC7947] Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server", RFC 7947, DOI 10.17487/RFC7947, September 2016, <<http://www.rfc-editor.org/info/rfc7947>>.

6.2. Informative 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>>.
- [RFC2622] Alaettinoglu, C., Villamizar, C., Gerich, E., Kessens, D., Meyer, D., Bates, T., Karrenberg, D., and M. Terpstra, "Routing Policy Specification Language (RPSL)", RFC 2622, DOI 10.17487/RFC2622, June 1999, <<http://www.rfc-editor.org/info/rfc2622>>.
- [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>>.
- [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>>.
- [RFC5291] Chen, E. and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", RFC 5291, DOI 10.17487/RFC5291, August 2008, <<http://www.rfc-editor.org/info/rfc5291>>.
- [RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <<http://www.rfc-editor.org/info/rfc5881>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <<http://www.rfc-editor.org/info/rfc6793>>.
- [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>>.
- [RS-ARCH] Govindan, R., Alaettinoglu, C., Varadhan, K., and D. Estrin, "A Route Server Architecture for Inter-Domain Routing", 1995, <<http://www.cs.usc.edu/assets/003/83191.pdf>>.

Acknowledgments

The authors would like to thank Chris Hall, Ryan Bickhart, Steven Bakker, and Eduardo Ascenco Reis for their valuable input.

Authors' Addresses

Nick Hilliard
INEX
4027 Kingswood Road
Dublin 24
Ireland

Email: nick@inex.ie

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

Email: elisa@bigwaveit.org

Robert Raszuk
Bloomberg LP
731 Lexington Ave.
New York, 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