

Network Working Group
Internet-Draft
Intended status: Informational
Expires: February 13, 2015

Pierre Francois
Institute IMDEA Networks
Camilo Cardona
Institute IMDEA Networks / UC3M
Adam Simpson
Alcatel-Lucent
Jeffrey Haas
Juniper Networks
August 12, 2014

ADD-PATH for Route Servers
draft-francois-idr-rs-addpaths-01

Abstract

BGP speakers at Internet Exchange Points typically exchange routes with a large number of peers. To reduce the burden of maintaining many sessions, IXPs implement and administrate BGP route servers. Route servers announce to their clients the paths of multiple peers by using a single eBGP session. Route servers, however, are restricted to propagating a single path per NLRI per eBGP session. This constraint affects the diversity of paths received by clients. To overcome this limitation, we propose the extension of ADD-PATH to eBGP peers. ADD-PATH allows a BGP speaker to send multiple paths for the same NLRI, typically through different nexthops, over a single session with a peer. By supporting ADD-PATH, a route server hence allows a client to potentially select among all the available paths for that NLRI, instead of the one arbitrarily chosen by the Route Server.

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 February 13, 2015.

Copyright Notice

Copyright (c) 2014 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 [2](#)
- [2.](#) Motivation [3](#)
- [3.](#) Operation of eBGP ADD-PATH capability for IXP route Server . 3
 - [3.1.](#) Capability [4](#)
 - [3.2.](#) Path Selection [4](#)
 - [3.2.1.](#) ADD-PATH ALL Policy compliant [4](#)
 - [3.2.2.](#) ADD-PATH N Policy compliant [4](#)
- [4.](#) Error conditions [5](#)
- [5.](#) IANA considerations [5](#)
- [6.](#) Security Considerations [5](#)
- [7.](#) References [5](#)
- Authors' Addresses [6](#)

1. Introduction

IXP route servers were designed to help network operators reduce the difficulties associated with maintaining a large number of sessions [[IXPRouteServer](#)]. Every route server client can receive paths from multiple ASes using the same eBGP session with the route server. In some cases, usually when there are many members in the IXP, multiple clients might announce a path to the same NLRI. Path diversity is an advantage for IXPs, as members can choose the path that better suits their policy. However, as a normal eBGP speaker, route servers can only advertise a single path per NLRI to each client. This limitation causes the route server to potentially hide paths that would be useful for their clients.

ADD-PATH [[AddPath](#)] is a capability that allows BGP speakers to announce more than one path to their peers. Works related to ADD-PATH have focused on applications for iBGP deployments. We propose the use of ADD-PATH over eBGP sessions to overcome the problems

associated with the limit on the number of paths that route servers can announce. In this document, we define the operation and error conditions of ADD-PATH for these scenarios and describe additional benefits for the route servers that implement it.

2. Motivation

By collecting paths from all their clients, route servers potentially accumulate various paths for some destination prefix. Multiple of these paths may be compliant with the policy of some clients of the route server. However, route servers typically maintain a single session with their clients, and hence advertise at most a single path towards each of them. As a result, a route server client will typically know only one of these paths. In some cases, depending on the implementation of the IXP route server, the route server client would not even receive a path for the NLRI. Refer to [\[IXPRouteServer\]](#) for an explanation of this problem.

We believe that this aspect of route serving is an unfortunate limitation, as it artificially hides paths from clients that may have wanted to use them.

First, it prevents the member from performing a policy based decision that is finer than the one advertised to the route server platform. That is, the arbitrary best path picked among the policy-compliant ones by the route server may be actually different from the one that the client would have picked, had it known about all of them.

Second, it prevents the member from doing temporary preference tweaking among the set of available paths in order to perform traffic engineering. That is, a member may only receive a path for a destination through a peer that is saturated, while alternate paths through non-saturated nexthops are available and would have been used if the router (and the operator) were aware of their existence.

ADD-PATH was designed to advertise more than one path towards a given NLRI. Multiple paths installed in the forwarding planes, as well as alternate paths, can be advertised among speakers supporting ADD-PATH. ADD-PATH can be used by a route server to announce all paths available for the same NLRI that still fulfill the policy of the route server client.

3. Operation of eBGP ADD-PATH capability for IXP route Server

3.1. Capability

A route server that supports the advertisement of multiple paths toward the same NLRI SHOULD announce the ADD-PATH capability to its clients. Likewise, a client supporting the reception of multiple paths SHOULD announce the ADD-PATH capability to the route server.

In an IXP context, only the route server should propagate multiple paths to the route server clients. The advertisement of multiple paths in the other direction is currently out of the specification of this document. Therefore, a route Server client should set the Send/Receive field for the Add-Path capability with a value of 1. The route Server should set the same field in the capability with a 2.

3.2. Path Selection

We describe here two path selection modes that can be implemented by the route server.

3.2.1. ADD-PATH ALL Policy compliant

A route server supporting ADD-PATH can announce to its clients all paths that comply with their policy. This selection mode is denominated as "ADD-PATH ALL Policy compliant".

3.2.2. ADD-PATH N Policy compliant

A route server may also support another type of ADD-PATH mode that restricts the number of paths per NLRI announced to each client. For instance, the route server would announce at most N paths to their clients that comply with their policies. This mode would help reduce the resources needed in the client, in case the number of available paths is large. Note that once the number of policy-compliant paths that can be advertised is restricted, a client might not receive the best possible path with respect to its own policies.

The configuration of the number of paths sent to each route server client could be done manually or set by the route server client via a communication channel.

The selection of paths is free to the implementation of the route server. Similarly to the ADD-PATH N mode [[AddPathGuidelines](#)], the route server COULD choose a set of paths equivalent to the one obtained after running the BGP best algorithm N times, excluding the selected path after each interaction.

4. Error conditions

In the specific context of route servers, third party nexthops are being used to have the client actually be able to select the appropriate nexthop. This is achieved by letting the route server leave the nexthop field of the propagated paths unchanged.

Similarly, the propagation of multiple paths by the route server to one of its clients must be made in a way that allows the receiver to actually select one among those paths. As a result, a route server advertising two different paths for the same destination, with equal nexthops, is out of specification. If this situation occurs, the client SHOULD log the event and let the normal decision process decide the best path.

As described in [Section 3.1](#), the advertisement of multiple paths in an IXP environment SHOULD only be performed from the route server to its clients. A route server SHOULD log any case in which a route server client signals, through the ADD-PATH capability, its willingness to announce more than one path. The route server SHOULD continue to operate under these events, considering all paths received from its clients.

5. IANA considerations

None

6. Security Considerations

The use of eBGP ADD-PATH in the route server environment does not increase the number of destinations for which paths are being advertised. However, the potential number of paths per destination is now larger than one, potentially increasing the memory load of the Adj-Rib-In. Systems risking to be short on memory due to this increase should be configured to constrain the amount of paths being advertised to them by a value which ensures proper operations.

7. References

[AddPath] Walton, D., Chen, E., Retana, A., and J. Scudder, "Advertisement of Multiple Paths in BGP", [draft-ietf-idr-add-paths-09.txt](#) (work in progress), October 2013.

[AddPathGuidelines]

Uttaro, J., Francois, P., Fragassi, R., Simpson, A., Patel, K., and P. Mohapatra, "Best Practices for Advertisement of Multiple Paths in IBGP", [draft-ietf-idr-add-paths-guidelines-06.txt](#) (work in progress), January 2014.

[IXPRouteServer]

Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange Route Server", [draft-ietf-idr-ix-bgp-route-server-05](#) (work in progress), June 2014.

Authors' Addresses

Pierre Francois
Institute IMDEA Networks
Avda. del Mar Mediterraneo, 22
Leganes 28918
ES

Email: pierre.francois@imdea.org

Camilo Cardona
Institute IMDEA Networks / UC3M
Avda. del Mar Mediterraneo, 22
Leganes 28918
ES

Email: juancamilo.cardona@imdea.org

Adam Simpson
Alcatel-Lucent
600 March Road
Ontario K2K 2E6
CA

Email: adam.simpson@alcatel-lucent.com

Jeffrey Haas
Juniper Networks
1194 N. Mathilda Ave
Sunnyvale 94089
USA

Email: jhaas@juniper.net

