

Network Working Group
Internet-Draft
Expires: August 11, 2014

Pierre Francois
Camilo Cardona
Institute IMDEA Networks
Adam Simpson
Alcatel-Lucent
Jeffrey Haas
Juniper Networks
February 7, 2014

ADD-PATH limit capability
draft-francois-idr-addpath-limit-00

Abstract

In this draft, we propose a new capability that allows BGP speakers supporting ADD-PATH to announce a limit on the number of paths they want to receive from their peers.

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 August 11, 2014.

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

Internet-Draft

ADD-PATH limit capability

February 2014

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	The ADD-PATH Limit capability	3
3.	Operation	4
4.	References	5
	Authors' Addresses	5

1. Introduction

ADD-PATH is an extension to BGP that allows a BGP speaker to send and receive more than one path for the same NLRI [[AddPath](#)]. The ADD-PATH capability does not include any mechanism for restricting the number or type of paths that a peer can receive from others. Thus, a receiving BGP speaker has no control on the number of paths sent by its peer, which depends only on the mode the operator desires to implement in the transmitting side [[AddPathGuidelines](#)]. This restriction can make network operations more difficult in some situations:

- o In cases in which all network devices are managed by the same group, operators select the ADD-PATH mode that best fits the resources of each of their devices. Operators must configure manually each speaker with a mode that does not overload the resources of the other devices. The overhead of this procedure can be high in some networks, as this configuration must be performed at the session level. If a ADD-PATH router could signal a limit in the number of paths it wants to receive, operators could achieve the same resource control by performing a more simple configuration.
- o In cases in which devices are managed by different operators, such as in networks spanning large geographical regions or in Internet Exchange Points, operators must agree on the ADD-PATH mode to use between BGP speakers. If one of the devices has constraints on the number of paths it can receive, the other party must configure an ADD-PATH mode that does not overload the memory of other device. Under these circumstances, allowing the receiving side to limit the number of paths can ease the management process for all administration sides.

In this document, we propose a new capability that allows an ADD-PATH capable BGP speaker to set an explicit upperbound on the number of paths it wants to receive from its peer.

2. The ADD-PATH Limit capability

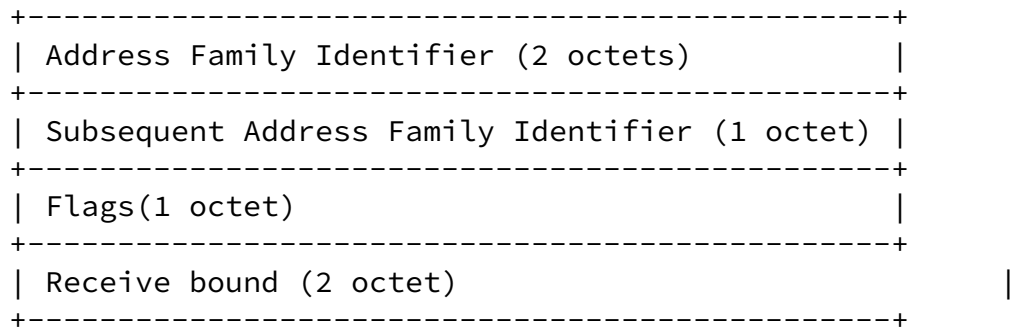


Figure 1: Capability

The meaning and use of the fields are as follows:

Address Family Identifier (AFI):

This field is the same as the one used in [\[RFC4760\]](#).

Subsequent Address Family Identifier (SAFI):

This field is the same as the one used in [\[RFC4760\]](#).

Flag Field

This field contains different bit flags related to the ADD-PATH limitation capability.

The most significant bit of the field (limit-capable bit) is used to communicate if the device supports the limitation of paths announced to the peer. If the router supports the ADD-PATH Limit capability, but it is not capable of limiting the number of announced paths, it should set the limit-capable bit to 0.

The remaining bits MUST be set to 0 and SHOULD be ignored upon receipt.

Receive Bound

This field indicates the maximum number of paths the device wants to receive from a peer for the <AFI, SAFI>. If the field is set to 0, the device has no restriction on the number of

paths it can receive.

[3.](#) Operation

The ADD-PATH Limit capability SHOULD be announced by BGP speakers that require their neighbors to send a limited number of paths per prefix. A router that is capable of sending a limited number of paths to a BGP ADD-PATH neighbor SHOULD also announce the ADD-PATH Limit capability. For cases in which a router is not capable of setting a limit in the number of paths it sends to a peer, it should set the limit-capable bit in the add-path capability to 0.

The ADD-PATH capability is a prerequisite of the ADD-PATH Limit. A BGP peer SHOULD ignore the ADD-PATH Limit capability from a peer that did not also announce the ADD-PATH capability.

The ADD-PATH limit capability is used to set an upper bound on the

number of paths that the router wants to receives from a peer. A sender capable of limiting the number of paths per peer SHOULD NOT send more paths than requested by the receiver.

It might be impractical for a BGP speaker to strictly stick to each of the upper bounds specified by its peers. Thus, the sender MAY send a lower amount of paths than the upper bound indicated by its peer.

A router SHOULD include the best path in the subset of paths to send to a peer, except when the best path is received from that peer. The choice of the rest of paths to be sent is left free to the implementation.

A BGP speaker could receive more paths than the number defined in the ADD-PATH capability, even when the BGP peer supports the limitation of paths. This event should be logged, but the session with the peer should be preserved. The receiving speaker should implement the required mechanism to deal with more paths that it can support.

[4.](#) References

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

[AddPathGuidelines]
J. Uttaro, P. Francois, R. Fragassi, A. Simpson, and K.
Patel, "Best Practices for Advertisement of Multiple Paths
in IBGP", [draft-ietf-idr-add-paths-guidelines-05.txt](#) (work
in progress).

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
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