

Workgroup: Inter-Domain Routing
Internet-Draft:
draft-abraitis-idr-addpath-paths-limit-00
Published: 17 January 2024
Intended Status: Standards Track
Expires: 20 July 2024
Authors: D. Abraitis A. Retana
 NetDef Futurewei Technologies, Inc.
 J. Haas
 Juniper Networks

Paths Limit for Multiple Paths in BGP

Abstract

This document specifies a BGP capability that complements the ADD-PATH Capability by indicating the maximum number of paths a BGP speaker can receive from a peer, optimizing the transmission of BGP routes by selectively relaying pertinent routes instead of the entire set.

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 <https://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 20 July 2024.

Copyright Notice

Copyright (c) 2024 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 (<https://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 Revised BSD License text as described in

Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Specification of Requirements](#)
- [3. PATHS-LIMIT Capability](#)
- [4. IANA Considerations](#)
- [5. Security Considerations](#)
- [6. References](#)
 - [6.1. Normative References](#)
 - [6.2. Informative References](#)
- [Acknowledgements](#)
- [Authors' Addresses](#)

1. Introduction

BGP ADD-PATH [[RFC7911](#)] defines a BGP extension that allows the advertisement of multiple paths for the same address prefix without the new paths implicitly replacing any previous ones.

Multiple paths for a large number of prefixes may be received by a BGP speaker, potentially depleting memory resources or even causing network-wide instability. Such instability could be considered a denial-of-service attack. Without knowing the maximum number of paths the receiver wants to receive, the sender may send more than that number of paths. [[I-D.ietf-idr-add-paths-guidelines](#)] provides recommendations for the use of BGP ADD-PATH while implementing specific applications.

This document specifies a BGP capability [[RFC5492](#)] that complements the ADD-PATH Capability [[RFC7911](#)] by indicating the maximum number of paths a BGP speaker can receive from a peer. This indication allows the sender to optimize the transmission of BGP routes by selectively relaying pertinent routes instead of the entire set.

2. Specification of Requirements

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 BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

3. PATHS-LIMIT Capability

The PATHS-LIMIT Capability is a BGP capability [[RFC5492](#)], with Capability Code TBD. The Capability Length field of this capability

is variable. The Capability Value field consists of one or more of the following tuples:




Figure 1

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 [\[RFC4760\]](#).

Paths Limit:

This field indicates the maximum paths limit the receiver wants to receive from its peer. If the received Paths Limit is zero (0), the tuple SHOULD be ignored.

A BGP speaker that wishes to indicate support for multiple AFI/SAFIs MUST do so by including the information in a single instance of the PATHS-LIMIT capability.

The PATHS-LIMIT capability MUST be ignored if the ADD-PATH capability is not present.

If the PATHS-LIMIT capability is empty (i.e. the Capability Length field is set to 0), it means that the sender doesn't have any specific limits to communicate.

An AFI/SAFI tuple MUST be ignored if the same tuple was not received in the ADD-PATH capability.

If more than one tuple is received for the same AFI/SAFI pair, only the first tuple should be considered. All others MUST be ignored.

A sender advertising multiple paths for the same prefix SHOULD send only the specified maximum number of paths indicated in the PATHS-LIMIT capability.

An implementation SHOULD provide a configuration knob to specify the maximum number of paths to accept from a sender.

4. IANA Considerations

IANA is requested to assign a value (TBD) for the PATHS-LIMIT Capability from the "Capability Codes" registry.

5. Security Considerations

This document defines a BGP extension that allows a receiver to better control the number of routes it receives when using BGP ADD-PATH [RFC7911]. Use of the PATHS-LIMIT capability can then mitigate some of the security-related concerns expressed in [RFC7911].

A rogue node or misconfiguration can result in the advertisement of a Paths Limit value that is too low for the application being used. This can result in inconsistent forwarding. Describing applications for BGP ADD-PATH is outside the scope of this document. Users of the PATHS-LIMIT Capability are encouraged to examine the behavior and potential impact by studying the best practices described in [I-D.ietf-idr-add-paths-guidelines].

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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <<https://www.rfc-editor.org/info/rfc5492>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

6.2. Informative References

[I-D.ietf-idr-add-paths-guidelines]

Uttaro, J., Francois, P., Patel, K., Haas, J., Simpson, A., and R. Fragassi, "Best Practices for Advertisement of Multiple Paths in IBGP", Work in Progress, Internet-Draft, draft-ietf-idr-add-paths-guidelines-08, 25 April 2016, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-add-paths-guidelines-08>>.

Acknowledgements

TBD

Authors' Addresses

Donatas Abraitis
NetDef

Email: donatas@opensourcerouting.org

Alvaro Retana
Futurewei Technologies, Inc.

Email: aretana@futurewei.com

Jeffrey Haas
Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089
United States of America

Email: jhaas@juniper.net