

Inter-Domain Routing
Internet-Draft
Updates: [4271](#) (if approved)
Intended status: Standards Track
Expires: 23 July 2022

M. Aelmans
Juniper Networks
M. Stucchi
Independent
J. Snijders
Fastly
19 January 2022

BGP Maximum Prefix Limits Inbound
draft-sas-idr-maxprefix-inbound-04

Abstract

This document describes mechanisms to limit the negative impact of route leaks [[RFC7908](#)] and/or resource exhaustion in BGP [[RFC4271](#)] implementations.

Requirements Language

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.

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 23 July 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

BGP Maximum Prefix Limits Inbound

January 2022

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
2.	Changes to RFC4271 Section 6	2
3.	Changes to RFC4271 Section 8	3
4.	Changes to RFC4271 Section 9	4
5.	Security Considerations	5
6.	IANA Considerations	5
7.	Acknowledgments	5
8.	Implementation status - RFC EDITOR: REMOVE BEFORE PUBLICATION	5
9.	Appendix: Implementation Guidance	6
10.	References	6
10.1.	Normative References	6
10.2.	Informative References	7
	Authors' Addresses	7

[1.](#) Introduction

This document updates [[RFC4271](#)] by revising control mechanism which limit the negative impact of route leaks [[RFC7908](#)] and/or resource exhaustion in Border Gateway Protocol (BGP) implementations. While [[RFC4271](#)] described methods to tear down BGP sessions or discard UPDATES after certain thresholds are exceeded, some nuances in this specification were missing resulting in inconsistencies between BGP implementations.

[2.](#) Changes to [RFC4271 Section 6](#)

This section updates [[RFC4271](#)] to specify what events can result in AutomaticStop (Event 8) in the BGP FSM.

The following paragraph replaces the second paragraph of [Section 6.7](#) (Cease), which starts with "A BGP speaker MAY support" and ends with

"The speaker MAY also log this locally.":

A BGP speaker MAY support the ability to impose a locally-configured, upper bound on the number of address prefixes the speaker is willing to accept from a neighbor (inbound maximum

prefix limit). The limit on the prefixes accepted from a neighbor can be applied before policy processing (Pre-Policy) or after policy processing (Post-Policy). When the upper bound is reached, the speaker, under control of local configuration, either:

- a. Discards new address prefixes from the neighbor, while maintaining the BGP connection. As these prefixes are discarded, their reachability information is not stored on the local router, which might lead to inconsistent routing behaviour;
- b. Receives all the new prefixes exceeding the threshold, accepts them and generates a log of the event;
- c. Terminates the BGP connection with the neighbor.

If the BGP speaker decides to terminate its BGP connection with a neighbor because the number of address prefixes received from the neighbor exceeds the locally-configured, upper bound, then the speaker MUST send the neighbor a NOTIFICATION message with the Error Code Cease.

Subcode	Symbolic Name
1	Threshold exceeded: Maximum Number of Prefixes Received

Table 1

The speaker MAY also log this locally.

3. Changes to [RFC4271 Section 8](#)

This section updates [Section 8 \[RFC4271\]](#), the paragraph that starts

with "One reason for an AutomaticStop event is" and ends with "The local system automatically disconnects the peer." is replaced with:

Possible reasons for an AutomaticStop event are: A BGP speaker receives an UPDATE messages with a number of prefixes for a given peer such that the total prefixes received exceeds the maximum number of prefixes configured (either "Pre-Policy" or "Post-Policy"). The local system automatically disconnects the peer.

4. Changes to [RFC4271 Section 9](#)

This section updates [\[RFC4271\]](#) by adding a subsection after [Section 9.4](#) (Originating BGP routes) to specify various events that can lead up to AutomaticStop (Event 8) in the BGP FSM.

9.5 Maximum Prefix Limits

9.5.1 Pre-Policy Inbound Maximum Prefix Limits

The Adj-RIB-In stores routing information learned from inbound UPDATE messages that were received from another BGP speaker [Section 3.2 \[RFC4271\]](#). The pre-policy limit uses the number of NLRIs per Address Family Identifier (AFI) per Subsequent Address Family Identifier (SAFI) as input into its threshold comparisons. For example, when an operator configures the pre-policy limit for IPv4 Unicast to be 50 on a given EBGP session, and the other BGP speaker announces its 51st IPv4 Unicast NLRI, the session MUST be terminated.

Pre-policy limits are particularly useful to help dampen the effects of full table route leaks and memory exhaustion when the implementation stores rejected routes.

Operators SHOULD take special care when utilizing methods where the router maintains a table of all the received updates pre-policy, as this could still expose control plane to exhaustion if no pre-policy limits are available or are not configured. Implementations SHOULD provide means to configure two

thresholds for inbound limits, one before policies are applied, and one after. This is to prevent exhaustion of control plane resources. The threshold before policy SHOULD be higher than or equal to the limit configured after policy.

9.5.2 Post-Policy Inbound Maximum Prefix Limits

[RFC4271](#) describes a Policy Information Base (PIB) that contains local policies that can be applied to the information in the Routing Information Base (RIB). The post-policy limit uses the number of NLRI per Address Family Identifier (AFI) per Subsequent Address Family Identifier (SAFI), after application of the Import Policy as input into its threshold comparisons. For example, when an operator configures the post-policy limit for IPv4 Unicast to be 50 on a given EBGP session, and the other BGP speaker announces a hundred IPv4 Unicast routes of which none are accepted as a result of the local import policy (and thus not considered for the Loc-RIB by the local BGP speaker), the session is not terminated.

Post-policy limits are useful to help prevent FIB exhaustion and prevent accidental BGP session teardown due to prefixes not accepted by policy anyway.

[5.](#) Security Considerations

Maximum Prefix Limits are an essential tool for routing operations and SHOULD be used to increase stability for the global routing ecosystem.

[6.](#) IANA Considerations

This memo requests that IANA updates the name of subcode "Maximum Number of Prefixes Reached" to "Threshold exceeded: Maximum Number of Prefixes Received" in the "Cease NOTIFICATION message subcodes" registry under the "Border Gateway Protocol (BGP) Parameters" group.

[7.](#) Acknowledgments

The authors would like to thank Saku Ytti and John Heasley (NTT Ltd.), Jeff Haas, Colby Barth and John Scudder (Juniper Networks), Martijn Schmidt (i3D.net), Teun Vink (BIT), Sabri Berisha (eBay),

Martin Pels (Quanza), Steven Bakker (AMS-IX), Aftab Siddiqui (ISOC), Yu Tianpeng, Ruediger Volk (Deutsche Telekom), Robert Raszuk (NTT), Jakob Heitz (Cisco), Warren Kumari (Google), Ben Maddison (Workonline), Randy Bush, Brian Dickson, Gyan Mishra (Verizon) and John John Heasley (NTTA) for their support, insightful reviews, and comments.

8. Implementation status - RFC EDITOR: REMOVE BEFORE PUBLICATION

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC7942](#). The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

The below table provides an overview (as of the moment of writing) of which vendors have produced implementation of inbound prefix limits. Each table cell shows the applicable configuration keywords if the vendor implemented the feature.

Vendor	Type A Pre-Policy	Type B Post-Policy
Cisco IOS XR		maximum-prefix
Cisco IOS XE		maximum-prefix
Juniper Junos OS	prefix-limit	accepted-prefix-limit, or prefix-limit combined with 'keep none'
Nokia SR OS	prefix-limit	

NIC.CZ BIRD	'import keep filtered' combined with 'receive limit'	'import limit' or 'receive limit'
OpenBSD OpenBGPD	max-prefix	
Arista EOS	maximum-routes	maximum-accepted-routes
Huawei VRPv5	peer route-limit	
Huawei VRPv8	peer route-limit	peer route-limit accept- prefix

Table 2: Maximum prefix limits capabilities per implementation

First presented by Snijders at [[RIPE77](#)]

[9.](#) Appendix: Implementation Guidance

TBD

[10.](#) References

[10.1.](#) 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>>.

[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,

<<https://www.rfc-editor.org/info/rfc4271>>.

[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>>.

10.2. Informative References

[RFC7908] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", [RFC 7908](#), DOI 10.17487/RFC7908, June 2016, <<https://www.rfc-editor.org/info/rfc7908>>.

[RIPE77] Snijders, J., "Robust Routing Policy Architecture", May 2018, <https://ripe77.ripe.net/wp-content/uploads/presentations/59-RIPE77_Snijders_Routing_Policy_Architecture.pdf>.

Authors' Addresses

Melchior Aelmans
Juniper Networks
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Netherlands

Email: maelmans@juniper.net

Massimiliano Stucchi
Independent

Email: max@stucchi.ch

Job Snijders
Fastly
Theodorus Majofskistraat 100
1065 SZ Amsterdam
Netherlands

