

Network Working Group
Internet-Draft
Updates: [4271](#) (if approved)
Intended status: Standards Track
Expires: August 15, 2017

R. Bush
Internet Initiative Japan
K. Patel
Arrcus, Inc.
D. Ward
Cisco Systems
February 11, 2017

Extended Message support for BGP
draft-ietf-idr-bgp-extended-messages-16

Abstract

The BGP specification mandates a maximum BGP message size of 4096 octets. As BGP is extended to support newer AFI/SAFIs, there is a need to extend the maximum message size beyond 4096 octets. This document updates the BGP specification [RFC 4271](#) by providing an extension to BGP to extend its current message size from 4096 octets to 65535 octets.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

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 15, 2017.

Copyright Notice

Copyright (c) 2017 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.	BGP Extended Message	2
3.	Extended message Capability for BGP	3
4.	Operation	3
5.	Error Handling	3
6.	Acknowledgements	4
7.	IANA Considerations	4
8.	Security Considerations	4
9.	References	4
9.1.	Normative References	4
9.2.	Informative References	4
	Authors' Addresses	5

[1.](#) Introduction

The BGP specification [RFC4271](#) [[RFC4271](#)] mandates a maximum BGP message size of 4096 octets. As BGP is extended to support newer AFI/SAFIs and newer capabilities (e.g., [[I-D.ietf-sidr-bgpsec-overview](#)]), there is a need to extend the maximum message size beyond 4096 octets. This draft provides an extension to BGP to extend its current message size limit from 4096 octets to 65535 octets.

[2.](#) BGP Extended Message

A BGP message over 4096 octets in length is a BGP Extended Message.

BGP Extended Messages have maximum message size of 65535 octets. The smallest message that may be sent consists of a BGP header without a data portion (19 octets).

Multi-octet fields MUST be in network byte order.

3. Extended message Capability for BGP

To advertise the BGP Extended Message Capability to a peer, a BGP speaker uses BGP Capabilities Advertisement [[RFC5492](#)]. By advertising the BGP Extended Message Capability to a peer, a BGP speaker conveys that it is able to send, receive, and properly handle BGP Extended Messages.

A peer which does not advertise this capability MUST NOT send BGP Extended Messages, and BGP Extended Messages MUST NOT be sent to it.

The BGP Extended Message Capability is a new BGP Capability [[RFC5492](#)] defined with Capability code TBD and Capability length 0.

4. Operation

A BGP speaker that is willing to send and receive BGP Extended Messages from its peer should advertise the BGP Extended Message Capability to its peer using BGP Capabilities Advertisement [[RFC5492](#)]. A BGP speaker may send extended messages to its peer only if it has received the Extended Message Capability from its peer.

An implementation that supports the BGP Extended Messages MUST be prepared to receive an OPEN message that is larger than 4096 bytes.

Applications generating messages which might be encapsulated within BGP messages MUST limit the size of their payload to take into account the maximum message size and all encapsulation overheads on the path the encapsulated data are expected to traverse.

5. Error Handling

A BGP speaker that has the ability to use extended messages but has not advertised the BGP Extended Messages capability, presumably due to configuration, SHOULD NOT accept an extended message. A speaker MAY implement a more liberal policy and accept extended messages even from a peer that has not advertised the capability.

However, a BGP speaker that does not advertise the BGP Extended Messages capability might also genuinely not support extended messages. Such a speaker would be expected to follow the error handling procedures of [[RFC4271](#)], [Section 6.1](#), and reset the session with a Bad Message Length NOTIFICATION if it receives an extended message. A speaker that treats an improper extended message as a fatal error, as described in the preceding paragraph, MUST do likewise.

The inconsistency between the local and remote BGP speakers MUST be reported via syslog and/or SNMP.

6. Acknowledgements

The authors thank Enke Chen, Susan Hares, John Scudder, John Levine, and Job Snijders for their input.

7. IANA Considerations

The IANA is requested to register a new BGP Capability Code to be named BGP Extended Message Capability and referring to this document.

Registry: BGP Capability Code

Value	Description	Document
-----	-----	-----
64	Graceful Restart Capability	[RFC4724]
....		
72	CP-ORF Capability	[RFC7543]
...		
TBD	BGP-Extended Message	[this draft]

8. Security Considerations

This extension to BGP does not change BGP's underlying security issues. It does enable large BGPsec BGPSEC_PATHs, see [I-D.ietf-sidr-bgpsec-protocol]

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), February 2009.

9.2. Informative References

- [I-D.ietf-sidr-bgpsec-overview]
Lepinski, M. and S. Turner, "An Overview of BGPSEC",
[draft-ietf-sidr-bgpsec-overview-02](#) (work in progress), May 2012.

[I-D.ietf-sidr-bgpsec-protocol]

Lepinski, M., "BGPSEC Protocol Specification", [draft-ietf-sidr-bgpsec-protocol-07](#) (work in progress), February 2013.

Authors' Addresses

Randy Bush
Internet Initiative Japan
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

Email: randy@psg.com

Keyur Patel
Arrcus, Inc.

Email: keyur@arrcus.com

Dave Ward
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: dward@cisco.com