IDR
Internet-Draft
Updates: 4486 (if approved)
Intended status: Standards Track
Expires: December 17, 2017

J. Snijders
NTT
J. Heitz
Cisco
J. Scudder
Juniper

June 15, 2017

BGP Administrative Shutdown Communication
draft-ietf-idr-shutdown-10

Abstract

This document enhances the BGP Cease NOTIFICATION message "Administrative Shutdown" and "Administrative Reset" subcodes for operators to transmit a short freeform message to describe why a BGP session was shutdown or reset. This document updates RFC 4486.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

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 December 17, 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

It can be troublesome for an operator to correlate a BGP-4 [RFC4271] session teardown in the network with a notice that was transmitted via off-line methods such email or telephone calls. This document updates [RFC4486] by specifying a mechanism to transmit a short freeform UTF-8 [RFC3629] message as part of a Cease NOTIFICATION message [RFC4271] to inform the peer why the BGP session is being shutdown or reset.

2. Shutdown Communication

If a BGP speaker decides to terminate its session with a BGP neighbor, and it sends a NOTIFICATION message with the Error Code "Cease" and Error Subcode "Administrative Shutdown" or "Administrative Reset" [RFC4486], it MAY include an UTF-8 encoded string. The contents of the string are at the operator's discretion.

The Cease NOTIFICATION message with a Shutdown Communication is encoded as below:

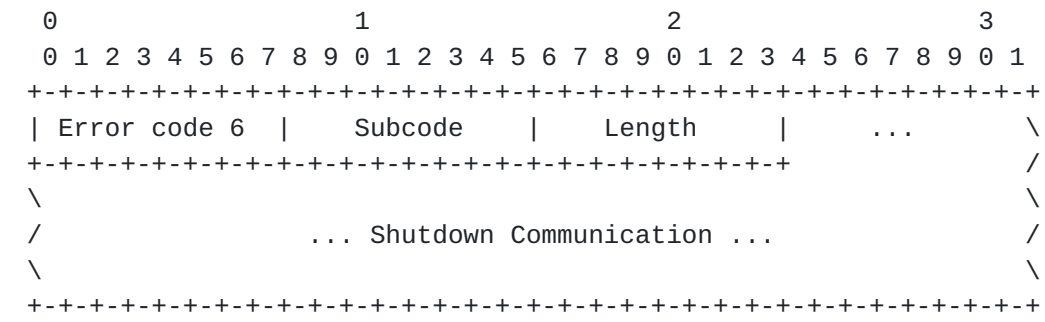

Figure 1

Subcode: the Error Subcode value MUST be one of the following values: 2 ("Administrative Shutdown") or 4 ("Administrative Reset").

Length: this 8-bit field represents the length of the Shutdown Communication field in octets. The length value MUST range from 0 to 128 inclusive. When the length value is zero, no Shutdown Communication field follows.

Shutdown Communication: to support international characters, the Shutdown Communication field MUST be encoded using UTF-8. A receiving BGP speaker MUST NOT interpret invalid UTF-8 sequences. Note that when the Shutdown Communication contains multibyte characters, the number of characters will be less than the length value. This field is not NUL terminated.

Mechanisms concerning the reporting of information contained in the Shutdown Communication are implementation specific but SHOULD include methods such as SYSLOG [RFC5424].

3. Operational Considerations

Operators are encouraged to use the Shutdown Communication to inform their peers of the reason for the shutdown of the BGP session and include out-of-band reference materials. An example of a useful Shutdown Communication would be:

"[TICKET-1-1438367390] software upgrade, back in 2 hours"

"[TICKET-1-1438367390]" is a ticket reference with significance to both the sender and receiver, followed by a brief human readable message regarding the reason for the BGP session shutdown followed by an indication about the length of the maintenance. The receiver can now use the string 'TICKET-1-1438367390' to search in their email archive to find more details.

4. Error Handling

If a Shutdown Communication with an invalid Length value, or an invalid UTF-8 sequence is received, a message indicating this event SHOULD be logged for the attention of the operator. An erroneous or malformed Shutdown Communication itself MAY be logged in a hexdump format.

5. IANA Considerations

Per this document, IANA is requested to reference this document at subcode "Administrative Shutdown", and at subcode "Administrative Reset" in the "Cease NOTIFICATION message subcodes" registry under the "Border Gateway Protocol (BGP) Parameters" group in addition to [RFC4486].

6. Security Considerations

This document uses UTF-8 encoding for the Shutdown Communication. There are a number of security issues with UNICODE. Implementers and operator are advised to review UNICODE TR36 [UTR36] to learn about these issues. UTF-8 "Shortest Form" encoding is REQUIRED to guard against the technical issues outlined in UTR36.

As BGP Shutdown Communications are likely to appear in syslog output, there is a risk that carefully constructed Shutdown Communication might be formatted by receiving systems in a way to make them appear as additional syslog messages. To limit the ability to mount such an attack, the BGP Shutdown Communication is limited to 128 octets in length.

Users of this mechanism should be aware that unless a transport that provides integrity is used for the BGP session in question, a Shutdown Communication message could be forged. Unless a transport that provides confidentiality is used, a Shutdown Communication message could be snooped by an attacker. These issues are common to any BGP message but may be of greater interest in the context of this proposal since the information carried in the message is generally expected to be used for human-to-human communication. Refer to the related considerations in [RFC4271] and [RFC4272].

Users of this mechanism should consider applying data minimization practises as outlined in Section 6.1 [RFC6973] as a received Shutdown Communication may be used at the receiver's discretion.

7. 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.

As of today these vendors have produced an implementation of the Shutdown Communication:

- ExaBGP
- pmacct
- OpenBGPD
- GoBGP
- FreeRangeRouting (frr)
- tcpdump (packet analyser)
- Wireshark (packet analyser)

8. References

8.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <http://www.rfc-editor.org/info/rfc3629>.

[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, <http://www.rfc-editor.org/info/rfc4271>.

[RFC4486] Chen, E. and V. Gillet, "Subcodes for BGP Cease Notification Message", RFC 4486, DOI 10.17487/RFC4486, April 2006, <http://www.rfc-editor.org/info/rfc4486>.

8.2. Informative References

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <http://www.rfc-editor.org/info/rfc4272>.

[RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, DOI 10.17487/RFC5424, March 2009, <http://www.rfc-editor.org/info/rfc5424>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <http://www.rfc-editor.org/info/rfc6973>.

[UTR36] Davis, M. and M. Suignard, "Unicode Security Considerations", Unicode Technical Report #36, August 2010, <http://unicode.org/reports/tr36/>.

Appendix A. Acknowledgements

The authors would like to gratefully acknowledge Tom Scholl, David Freedman, Jared Mauch, Jeff Haas, Peter Hessler, Bruno Decraene, John Heasley, Peter van Dijk, Arjen Zonneveld, James Bensley, Susan Hares, Saku Ytti, Lou Berger, Alvaro Retana, and Adam Roach.

The authors would like to thank Enke Chen and Vincent Gillet for their work on [RFC4486] and granting the related BCP 78 rights to the IETF Trust.

Authors' Addresses

Job Snijders
NTT Communications
Theodorus Majofskistraat 100
Amsterdam 1065 SZ
The Netherlands

Email: job@ntt.net

Jakob Heitz
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

Email: jheitz@cisco.com

John Scudder
Juniper Networks
1194 N. Mathilda Ave
Sunnyvale, CA 94089
USA

Email: jgs@juniper.net