

Subcodes for BGP Cease Notification Message

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document defines several subcodes for the BGP Cease NOTIFICATION message that would provide more information to aid network operators in correlating network events and diagnosing BGP peering issues.

1. Introduction

This document defines several subcodes for the BGP Cease NOTIFICATION message that would provide more information to aid network operators in correlating network events and diagnosing BGP peering issues. It also recommends that a BGP speaker implement a backoff mechanism in re-trying a BGP connection after the speaker receives a NOTIFICATION message with certain CEASE subcode.

2. Specification of Requirements

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 RFC 2119 [RFC-2119].

3. Subcode Definition

The following subcodes are defined for the Cease NOTIFICATION message:

Subcode	Symbolic Name
1	Maximum Number of Prefixes Reached
2	Administrative Shutdown
3	Peer De-configured
4	Administrative Reset
5	Connection Rejected
6	Other Configuration Change
7	Connection Collision Resolution
8	Out of Resources

4. Subcode Usage

If a BGP speaker decides to terminate its peering with a neighbor because the number of address prefixes received from the neighbor exceeds a locally configured upper bound (as described in [BGP-4]), then the speaker MUST send to the neighbor a NOTIFICATION message with the Error Code Cease and the Error Subcode "Maximum Number of Prefixes Reached". The message MAY optionally include the Address Family information [BGP-MP] and the upper bound in the "Data" field, as shown in Figure 1, where the meaning and use of the <AFI, SAFI> tuple is the same as defined in [BGP-MP], Section 7.




Figure 1: Optional Data Field

If a BGP speaker decides to administratively shut down its peering with a neighbor, then the speaker SHOULD send a NOTIFICATION message with the Error Code Cease and the Error Subcode "Administrative Shutdown".

If a BGP speaker decides to de-configure a peer, then the speaker SHOULD send a NOTIFICATION message with the Error Code Cease and the Error Subcode "Peer De-configured".

If a BGP speaker decides to administratively reset the peering with a neighbor, then the speaker SHOULD send a NOTIFICATION message with the Error Code Cease and the Error Subcode "Administrative Reset".

If a BGP speaker decides to disallow a BGP connection (e.g., the peer is not configured locally) after the speaker accepts a transport protocol connection, then the BGP speaker SHOULD send a NOTIFICATION message with the Error Code Cease and the Error Subcode "Connection Rejected".

If a BGP speaker decides to administratively reset the peering with a neighbor due to a configuration change other than the ones described above, then the speaker SHOULD send a NOTIFICATION message with the Error Code Cease and the Error Subcode "Other Configuration Change".

If a BGP speaker decides to send a NOTIFICATION message with the Error Code Cease as a result of the collision resolution procedure (as described in [BGP-4]), then the subcode SHOULD be set to "Connection Collision Resolution".

If a BGP speaker runs out of resources (e.g., memory) and decides to reset a session, then the speaker MAY send a NOTIFICATION message with the Error Code Cease and the Error Subcode "Out of Resources".

It is RECOMMENDED that a BGP speaker behave as though the DampPeerOscillations attribute [BGP-4] were true for this peer when re-trying a BGP connection after the speaker receives a Cease NOTIFICATION message with a subcode of "Administrative Shutdown", "Peer De-configured", "Connection Rejected", or "Out of Resources". An implementation SHOULD impose an upper bound on the number of consecutive automatic retries. Once this bound is reached, the implementation would stop re-trying any BGP connections until some administrative intervention, i.e., set the AllowAutomaticStart attribute [BGP-4] to FALSE.

5. IANA Considerations

This document defines the subcodes 1 - 8 for the BGP Cease NOTIFICATION message. Future assignments are to be made using either the Standards Action process defined in [RFC-2434], or the Early IANA Allocation process defined in [RFC-4020]. Assignments consist of a name and the value.

6. Security Considerations

This extension to BGP does not change the underlying security issues inherent in the existing BGP.

7. Acknowledgements

The authors would like to thank Yakov Rekhter, Pedro Marques, Andrew Lange, and Don Goodspeed for their review and suggestions.

8. References

8.1. Normative References

- [BGP-4] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [BGP-MP] Bates, T., Rekhter, Y., Chandra, R., and D. Katz, "Multiprotocol Extensions for BGP-4", RFC 2858, June 2000.
- [RFC-2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [RFC-2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

- [RFC-4020] Kompella, K. and A. Zinin, "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 4020, February 2005.

Authors' Addresses

Enke Chen
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134
USA

EMail: enkechen@cisco.com

Vincent Gillet
France Telecom Longues Distances
61, rue des Archives
75003 Paris FRANCE

EMail: vgi@opentransit.net

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

