

Network Working Group
Internet Draft
Expires: January 2006

Gargi Nalawade
Keyur Patel
John Scudder
David Ward

Cisco Systems

BGPv4 Soft-Notification Message

[draft-nalawade-bgp-soft-notify-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

This document defines a new message type, the BGP SOFT-NOTIFICATION message that can notify a peer of an error-condition for a particular AFI/SAFI and soft-reset the AFI/SAFI, without terminating the BGP session and without impacting the other AFI/SAFIs exchanged with the peer.

1. Introduction

Currently there is no mechanism available for two BGP Speakers to communicate the occurrence of an error-condition other than through a BGP NOTIFICATION Message. The problem is that a NOTIFICATION message resets the peering session and terminates the connection. If a peer wants to gracefully recover from an error or wants to warn its peer about the occurrence of a BGP-related event, there is no mechanism currently available to do that.

The proposed BGP SOFT-NOTIFICATION message is a mechanism to notify a remote peer of an error-condition or an event without resetting or terminating the BGP session. The purpose of this message is to provide the ability to soft-reset a particular AFI/SAFI without disrupting other BGP AFI/SAFIs or sections.

2. Definition of the BGP SOFT-NOTIFICATION Message

The BGP SOFT-NOTIFICATION message is a BGP message with type TBD. A BGP SOFT-NOTIFICATION message may be sent to notify a peer of an error condition within a particular AFI/SAFI.

In addition to the fixed-size BGP header, the SOFT-NOTIFICATION message contains the following fields:

```

      0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
+---+---+---+---+---+---+---+---+
|           AFI           |
+---+---+---+---+---+---+---+---+
|           SAFI           |
+---+---+---+---+---+---+---+---+
|           Type-code       |
+---+---+---+---+---+---+---+---+
|           Sub-code        |
+---+---+---+---+---+---+---+---+
|           Length          |
+---+---+---+---+---+---+---+---+
|           Variable Data TLV
|
+---+---+---+---+---+---+---+---+
```

[draft-nalawade-bgp-soft-notify-01.txt](#)

2]

[Page

where :

AFI/SAFI : as defined in [[IANA-AFI](#)] [[IANA-SAFI](#)]

Type-code : is a 2 octet field indicating the error-condition or event for the respective AFI/SAFI.

Sub-code : is a 2-octet field that may define a sub-code related to the error-condition or event that this message carries.

Length : is a 2-octet field that contains the length of the remaining message that may contain an optional Variable length Data TLV.

Variable Data TLV : is a variable-length field that contains a Variable Data TLV that may be used to carry additional data to provide more information about the error-condition or event.

[2.1 AFI-SAFI](#)

The AFI-SAFI fields indicate the AFI/SAFI for which the error-condition or event has occurred and needs to be soft-reset. A Reserved Value [TBD] will indicate that this message is for all AFI/SAFIs. A Reserved Value [TBD] in only the SAFI field will indicate that this message applies to all SAFIs under a particular AFI.

[2.2 Type-code](#)

The following Type-codes have been defined :

Error Code	Symbolic Name
1	UPDATE Message Error
2	Cease
3	Event

[2.3 Sub-code](#)

The following Sub-codes have been defined :

UPDATE Message Error subcodes:

- 1 - Malformed Attribute List.
- 2 - Unrecognized Well-known Attribute.
- 3 - Missing Well-known Attribute.
- 4 - Attribute Flags Error.

[draft-nalawade-bgp-soft-notify-01.txt](#)

3]

[Page

- 5 - Invalid Attribute Length.
- 6 - Invalid ORIGIN Attribute.
- 7 - Invalid NEXT_HOP Attribute.
- 8 - Optional Attribute Error.
- 9 - Invalid Network Field.
- 10 - Bad ASPATH.
- 11 - Invalid Update-Message Type.

CEASE Message Error subcodes:

- 1 - Maximum Number of Prefixes Reached
- 2 - Administratively Shutdown
- 3 - Peer Unconfigured
- 4 - Administratively Reset
- 5 - Other Configuration Change

Event Message subcodes :

- 1 - ACK Soft-Notification
- 2 - Peer Administratively unshut
- 3 - Peer Configured
- 4 - Timer Expired
- 5 - Dampening routes
- 6 - Undampened routes

2.4. Variable Data TLV

The Variable Data TLV MAY be used to carry additional information about the Error or Event. It is of the following form :

```
+-----+
| Type (1 octet)      |
+-----+
| Length (2 octets)   |
+-----+
| Value (variable)    |
+-----+
```

This document defines TLVs summarized below:

Type	Name	Length	Value
----	-----	-----	-----
1	String	variable	A text string whose length is given by the length field. Not

[draft-nalawade-bgp-soft-notify-01.txt](#)

4]

[Page

			null-terminated.
2	PDU	variable	A copy of the PDU which triggered the SOFT-NOTIFICATION message. May be truncated.
3	Attribute	variable	A copy of the path attribute which triggered the SOFT-NOTIFICATION message. May be truncated.
4	Integer	4 octets	A four-octet integer

No TLV may appear in the SOFT-NOTIFICATION message more than once.

3. Operation

A BGP Speaker may generate a SOFT-NOTIFICATION for the relevant AFI/SAFIs in lieu of the NOTIFICATION Message [[BGP-4](#)] [[BGP-CEASE](#)] for the relevant Type-codes and sub-codes in [[BGP-4](#)] [and BGP-CEASE], as redefined in [section 2.2](#) for SOFT-NOTIFICATIONS. A BGP Speaker may also generate a SOFT-NOTIFICATION in case of an Event that is listed in the Event sub-codes above.

The following rules apply to the processing of the BGP SOFT-NOTIFICATION Message.

When a peer receives a BGP SOFT-NOTIFICATION Message, it will take an action based on the Type-code contained in the message. The sending peer will also take an action after it has sent the SOFT-NOTIFICATION out to its peer.

In the sections below the BGP Speaker sending the SOFT-NOTIFICATION is referred to as the Sending Speaker and the BGP Speaker receiving the SOFT-NOTIFICATION is referred to as the Receiving Speaker. A SOFT-NOTIFICATION Message with Type-Code "Event" and sub-code "Soft-Notify-ACK" is referred to simply as the "Soft-Notify-ACK".

3.1 Update Error Type-code

The Sending Speaker, upon sending a SOFT-NOTIFICATION message, will start a timer for the receipt of a Soft-Notify-ACK, and will discard any updates received from the Receiving Speaker after this, until a Soft-Notify-ACK is received. The Sending Speaker will then flush the routes of the Receiving Speaker for that AFI/SAFI and start sending out new updates to the Receiving Speaker. When the Sending Speaker receives the Soft-Notify-ACK, it will resume accepting updates from the Receiving Speaker. If the Soft-Notification Timer expires before receiving the Soft-Notify-ACK, the Sending Speaker will terminate the

[draft-nalawade-bgp-soft-notify-01.txt](#)

5]

[Page

BGP session.

If the SOFT-NOTIFICATION message received by the Receiving Speaker, contains an Update Error Type-code, the Receiving Speaker must send a Soft-Notify-ACK back to the Sending Speaker. The Receiving Speaker will also flush the routes of the Sending Speaker for that AFI/SAFI and then proceed to re-advertise its own routes to the Sending Speaker.

The exact hand-shaking mechanism is described in the diagram below.



[draft-nalawade-bgp-soft-notify-01.txt](#)

6]

[Page

3.2 Cease Type-code

The Sending Speaker, upon sending a SOFT-NOTIFICATION message with a Type-code Cease, must flush the routes of the Receiving Speaker for that AFI/SAFI and put the AFI/SAFI in a shutdown state for that peer.

The Receiving Speaker, upon receiving a SOFT-NOTIFICATION message with a Type-code Cease, must send a Soft-Notify-ACK to the Sending Speaker, flush the routes of the Sending Speaker for that AFI/SAFI and put the AFI/SAFI in a shutdown state for that peer. A shutdown state of a SAFI for a peer is a state in which the BGP Speaker will not accept and process any updates from the peer.

3.3 Event Type-code

The Sending Speaker, upon sending a SOFT-NOTIFICATION message with a Type-Code "Event" and subcode "Administratively unshut", must transition out of the shutdown state for that AFI/SAFI for that peer. It will then readvertise its routes for that AFI/SAFI to the Receiving Speaker.

The Receiving Speaker, upon receiving a SOFT-NOTIFICATION message with Type-Code "Event" and subcode "Administratively unshut", must send a Soft-Notify-ACK to the Sending Speaker and transition out of the shutdown state for that AFI/SAFI. The Receiving Speaker will then re-advertise its routes for the relevant AFI/SAFI to the Sending Speaker.

If the SOFT-NOTIFICATION message received contains any Event Type-code other than "Administratively unshut" and "Soft-Notify-ACK", the Receiving Speaker must send a Soft-Notify-ACK to the Sending Speaker and MAY choose to log the message.

3.4 Multiple Soft-Notifications

The sending of SOFT-NOTIFICATION Messages and soft-reset of the peer for an AFI/SAFI, should be rate-limited and some mechanism for exponential backoff should be implemented. The exact mechanism to be used is beyond the scope of this document.

If a Sending Speaker sends multiple SOFT-NOTIFICATION Messages, it must remember how many such Messages have not yet been acknowledged. When it receives a Soft-Notify-ACK, it must associate it with the earliest SOFT-NOTIFICATION message pending a Soft-Notify-ACK.

4. Capability

[draft-nalawade-bgp-soft-notify-01.txt](#)

7]

[Page

A new capability [[BGP-CAP](#)] code (TBD) is defined for the BGP SOFT-NOTIFICATION messages. A BGP SOFT-NOTIFICATION message can only be sent to peers that have advertised this capability.

5. Deployment considerations

This draft does not affect the deployment considerations for BGP.

6. Intellectual Property Considerations

Cisco Systems may seek patent or other intellectual property protection for some of all of the technologies disclosed in this document. If any standards arising from this document are or become protected by one or more patents assigned to Cisco Systems, Cisco intends to disclose those patents and license them on reasonable and non-discriminatory terms.

7. Security Considerations

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

8. Acknowledgements

The authors would like to thank Nehal Bahu, Shyam Suri, David Ball, Srihari R. Sangli and Pedro Marquez for their review and comments.

9. Normative References

[IANA-AFI] <http://www.iana.org/assignments/address-family-numbers>

[IANA-SAFI] <http://www.iana.org/assignments/safi-namespaces>

[BGP-4] Rekhter, Y. and T. Li (editors), "A Border Gateway Protocol 4 (BGP-4)", Internet Draft [draft-ietf-idr-bgp4-21.txt](#), March 2004.

[BGP-CAP] Chandra, R., Scudder, J., "Capabilities Advertisement with BGP-4", [draft-ietf-idr-rfc2842bis-02.txt](#), April 2002.

[BGP-CEASE] Chen, E., Gillett V., "Subcodes for BGP Cease Notification Message", [draft-ietf-idr-cease-subcode-03.txt](#), May 2003.

[BGP-INFORM] Nalawade, G., Scudder, J., Ward, D., "BGPv4 Inform message", [draft-nalawade-bgp-inform-02.txt](#), August 2002.

10. Author's Addresses

Gargi Nalawade <mailto:gargi@cisco.com>

[draft-nalawade-bgp-soft-notify-01.txt](#)

8]

[Page

Keyur Patel <mailto:keyupate@cisco.com>

John Scudder <mailto:jgs@cisco.com>

David Ward <mailto:dward@cisco.com>

Cisco Systems, Inc 170 West Tasman Drive San Jose, CA 95134

11. Full Copyright Statement

Copyright (C) The Internet Society (2005). All Rights Reserved.

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.

12. Expiration Date

This memo is filed as <[draft-nalawade-bgp-soft-notify-01.txt](#)> expires January, 2006.

