

Internet Draft

Document: [draft-nalawade-bgp-inform-02.txt](#)

Expires: December 2002

Gargi Nalawade

John Scudder

David Ward

Cisco Systems

June 2000

BGPv4 INFORM message

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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/lid-abstracts.txt>

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

2. Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

3. Abstract

This document defines a new message type, the BGP INFORM message that communicates Informational data and operational warnings without resetting the peering session.

[4. Table of Contents](#)

1. Status of this Memo.....	1
2. Copyright Notice.....	1
3. Abstract.....	1
5. Introduction.....	3
6.....	3
Definition of the BGP INFORM Message.....	3
6.1. Event Codes.....	4
6.1.1. Unspecified event.....	4
6.1.2. Recoverable UPDATE attribute error -- attribute discarded.....	5
6.1.3. Recoverable UPDATE attribute error -- attribute fixed.....	5
6.1.4. Too many routes -- routes discarded.....	5
6.1.5. Attribute Overflow.....	5
6.1.6. Dampening routes.....	5
6.1.7. All routes undampened.....	6
6.1.8. Graceful Restart Purge Timer Expired.....	6
7. Operation.....	6
7.1.1. Sending an INFORM Message.....	6
7.1.2. Receiving an INFORM message.....	6
7.1.3. Implementation notes.....	7
7.1.4. Capability.....	7
8. Security Considerations.....	8
9. Acknowledgements.....	8
10. References.....	8
11. Author's Addresses.....	8
12. Intellectual Property Statement.....	9
13. Full Copyright Statement.....	9
14. Expiration Date.....	11

Nalawade Internet Draft - Expires December 2002

[5. Introduction](#)

Currently there is no mechanism available for two peers to communicate the occurrence of an event other than through a BGP NOTIFICATION Message. The problem is that a NOTIFICATION message resets the peering session. 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 available to do that. The proposed BGP INFORM message is a mechanism to inform a remote peer of an event without resetting the session.

6. Definition of the BGP INFORM Message

The INFORM message is a BGP message with type TBD. An INFORM message may be sent to inform a peer of an error condition which is not serious enough to warrant the reset of the BGP peering session. Each INFORM message relates to a single event. To inform a peer about multiple events, multiple INFORM messages must be used.

The INFORM message contains a 2-octet Event Code followed by one or more Data TLVs of the following form:

```

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

This document defines TLVs summarized below:

Type	Name	Length	Value
1	Unspecified	variable	Unspecified Data Type.
2	String	variable	A text string whose length is given by the length field. Not null-terminated.
3	PDU	variable	A copy of the PDU which triggered

Nalawade Internet Draft - Expires December 2002
BGPv4 INFORM Message

August 200

the INFORM message. May be truncated.

4	Attribute	variable	A copy of the path attribute which triggered the INFORM message. May be truncated.
5	Integer	4	A four-byte integer

No TLV may appear in the INFORM message more than once.

[6.1.](#) Event Codes

The Event Code provides structured information regarding the event which triggered the generation of the INFORM message.

Events 0-32767 are well-known and are defined here (TBD IANA document). Events 32768-65535 are reserved for vendor-specific use.

Well-known events are summarized in the table below, and subsequently described.

Code	Name
----	----
1	Unspecified event
2	Recoverable UPDATE attribute error -- attribute Discarded
3	Recoverable UPDATE attribute error -- attribute fixed
4	Too many routes -- routes discarded
5	Attribute Overflow
6	Dampening routes
7	All routes undampened
8	Graceful Restart purge timer expired

In the descriptions below, the inclusion of certain TLVs is specified -- for example, an unspecified event should include a string describing the event. These constitute a minimum set that should be included -- any other applicable or useful TLV may also be included.

[6.1.1.](#) Unspecified event

An event has occurred which is not described by any other event code. The String TLV should be included with a description of the event.

[6.1.2.](#) Recoverable UPDATE attribute error -- attribute discarded

The attribute which caused the INFORM to be generated should be included in the Attribute TLV. The reason it was considered an error should be included in the String field of the data port of the packet.

6.1.3. Recoverable UPDATE attribute error -- attribute fixed

The attribute which caused the INFORM to be generated should be included in the Attribute TLV. The reason it was considered an error and a description of the action taken to fix the problem should be included in the String field of the data port of the packet.

Care should be taken not to fix attributes unless it can be unambiguously determined that doing so will not compromise the protocol's correctness.

6.1.4. Too many routes -- routes discarded

The peer has sent more routes than the local BGP speaker's configured maximum. The local BGP speaker has discarded some of the routes received from the peer.

The configured maximum value which was exceeded should be included in the Integer TLV.

6.1.5 Attribute Overflow

This INFORM message may be sent any time a peer receives an UPDATE with an attribute value, e.g., community list [[RFC1997](#)] that must be truncated due to its length. The Attribute TLV should be included.

6.1.6 Dampening routes

The remote peer has announced and withdrawn some prefix or prefixes too frequently and the local peer has applied dampening to some set of prefixes announced by the remote peer. This INFORM message should not be sent each time a prefix is dampened. Instead it should only be sent when the boundary from no dampened routes to any dampened routes has been crossed.

6.1.7 All routes undampened

At some point in the past, the remote peer announced and withdrawn some prefix or prefixes too frequently and the local peer had applied dampening to some set of prefixes announced by the remote peer. This INFORM message should not be sent each time a prefix is undampened. Instead it should only be sent when the boundary from dampened routes to no dampened prefixes has been crossed.

6.1.8 Graceful Restart Purge Timer Expired

This INFORM message is to be sent during a Graceful Restart event [[BGP-GR](#)] and the purge timer has expired, thus causing all routes from the remote peer to be purged from the forwarding table of the local peer.

7. Operation

The following rules apply to the generation of INFORM messages:

7.1.1. Sending an INFORM Message

A router may send an INFORM message to a peer upon detecting a normal or abnormal, non-critical condition during operation which needs to be communicated to the peer and which does not necessitate a session reset.

A router SHOULD NOT send an INFORM message for a condition which requires a session reset. It SHOULD NOT be sent in conjunction with a NOTIFICATION message.

The rate at which INFORM messages are generated must be rate-limited. A suggested default limit is 60 messages per minute.

7.1.2. Receiving an INFORM message

On receiving an INFORM Message from a peer, the INFORM message should be logged and via locally determined means and brought to the attention of the router's operator. The means to do this are, however, outside the scope of this draft.

Under all circumstances an implementation SHOULD NOT take any automated action upon receiving an INFORM message (other than logging or alerting the operator).

An implementation must be prepared to receive INFORM messages containing unrecognized TLVs or TLV subcodes. An implementation should handle recognized TLVs as normal and may log, silently drop, or otherwise handle unrecognized TLVs. It is not recommended that the reception of a malformed INFORM message be cause to generate a reply of an INFORM message. An implementation must not reset the session due to a malformed INFORM message.

7.1.3. Implementation notes

An implementation must not assume that its generation of an INFORM message will result in any state change on the part of its peer. It is axiomatic that the INFORM message is for the peer's information only.

Implementors should refrain from sending INFORM messages without good cause. Although use of an INFORM message is not as serious as sending a NOTIFICATION, nonetheless an INFORM should only be generated in response to a protocol error or other serious problem. Normal, expected protocol events should not be INFORMed. Examples of events for which generation of an INFORM would be inappropriate include the dampening of an individual flapping route, the impending expiration of a holdtime, or the suppression of a component of an aggregate.

The INFORM should not be used as a blanket replacement for sending a notification and terminating the BGP session. The BGP protocol's correctness generally assumes that protocol errors will be handled by terminating the session. The decision not to terminate a session in response to an error condition should not be taken lightly or without careful and sober consideration. It is noted that it is possible for a NOTIFICATION message to carry arbitrary data, so if the session is to be terminated, any relevant data may be carried in the NOTIFICATION message itself.

7.1.4. Capability

A new Capability [[BGP-CAP](#)] code (TBD) is defined for interoperability with software that does not recognize the INFORM message. The INFORM message can be sent only to peers that have advertised this capability.

8. Security Considerations

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

9. Acknowledgements

We would like to thank Russ White, Alvaro Retana and Chandrashekhar Appanna for their review of the document. The authors would also like to thank all the other reviewers that gave suggestions to this document.

10. References

[BGP-4] Rekhter, Y. and T. Li (editors), "A Border Gateway Protocol 4 (BGP-4)", Internet Draft [draft-ietf-idr-bgp4-17.txt](#), January 2002.

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

[BGP-GR] Chandra, R., Scudder, J., "Graceful Restart Mechanism for BGP", [draft-ietf-idr-restart-05.txt](#), June 2002.

[RFC1997] Chandra, R., Traina, P., Li, T., "BGP Communities Attribute", [RFC1997](#), August 1996.

11. Author's Addresses

Gargi Nalawade
mailto:gargi@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

12. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards- related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

13. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English. The limited permissions

granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns. This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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."

Nalawade Internet Draft - Expires December 2002
BGPv4 INFORM Message

1
August 200

14. Expiration Date

This memo is filed as <[draft-nalawade-bgpv4-INFORM-00.txt](#)>, and expires December, 2002.

Nalawade

Internet Draft - Expires December 2002

1