

Dynamic Capability for BGP-4

[draft-ietf-idr-dynamic-cap-06.txt](#)

1. Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, or will be disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) except that the right to produce derivative works is not granted.

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>.

2. Abstract

This document defines a new BGP capability termed "Dynamic Capability", which would allow the dynamic update of capabilities over an established BGP session. This capability would facilitate non-disruptive capability changes by BGP speakers.

3. Introduction

Currently BGP capabilities [[BGP-CAP](#)] are only advertised in the OPEN message during the session initialization. In order to enable a new capability or remove an existing capability (such as an Address Family support [[BGP-MP](#)]), an established session needs to be reset, which may disrupt other services running over the session.

This document defines a new BGP capability termed "Dynamic Capability", which would allow the dynamic update of capabilities over an established BGP session. This capability would facilitate non-disruptive capability changes by BGP speakers.

4. 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](#)].

5. Dynamic Capability

The Dynamic Capability is a new BGP capability [[BGP-CAP](#)]. The Capability Code for this capability is specified in the "IANA Considerations" section of this document. The Capability Length field of this capability is one octet. The Capability Value field consists of a list of capability codes (one-octet for each) for which the dynamic revision is supported by a BGP speaker.

By advertising the Dynamic Capability to a peer in the OPEN, a BGP speaker conveys to the peer that the speaker is capable of receiving and properly handling the CAPABILITY message (as defined in the next Section) from the peer after the BGP session has been established.

6. Capability Message

The CAPABILITY Message is a new BGP message type with type code 6. In addition to the fixed-size BGP header [[BGP-4](#)], the CAPABILITY message contains one or more of the following tuples of capability revisions:

```
+-----+
| Init/Ack (1 bit)          |
+-----+
| Ack Request (1 bit)      |
+-----+
| Reserved (5 bits)        |
+-----+
| Action (1 bit)           |
+-----+
| Sequence Number (4 octets) |
+-----+
| Capability Code (1 octet)  |
+-----+
| Capability Length (2 octets) |
+-----+
| Capability Value (variable) |
+-----+
```

The Init/Ack bit indicates whether a capability revision is being initiated (when set to 0), or being acknowledged (when set to 1).

The Ack Request bit indicates whether an acknowledgement is requested (when set to 1), or not (when set to 0) for a capability revision being initiated.

The Reserved bits should be set to zero by the sender and ignored by the receiver.

The Action bit is 0 for advertising a capability, and 1 for removing a capability.

The Sequence Number field can be used by a BGP speaker to match an acknowledgement with a capability revision that the speaker initiated previously.

Conceptually the triple <Capability Code, Capability Length, Capability Value> is the same as the one defined in [[BGP-CAP](#)], and it specifies a capability for which the "Action" shall be applied. The triple is optional when the Init/Ack bit is set to 1.

7. Operation

A BGP speaker that is willing to receive the CAPABILITY message (for one or more capability codes) from its peer SHOULD use the BGP Capabilities Advertisement [[BGP-CAP](#)] to advertise the Dynamic Capability for these capability codes.

A BGP speaker MAY send to its peer a CAPABILITY message to initiate revisions for one or more capability codes only if these capability codes are listed in the Dynamic Capability of the OPEN message received from its peer.

A CAPABILITY message MAY be received only in the Established state. Receiving a CAPABILITY message in any other state is a Finite State Machine Error as defined in [[BGP-4](#)]. A BGP speaker SHOULD reset the HoldTimer upon receiving a CAPABILITY message from its peer.

When a BGP speaker sends a CAPABILITY message to its peer to initiate a capability revision, the Init/Ack bit for the capability revision in the message MUST be set to 0. The setting of the Ack Request bit is capability specific. The assignment of the Sequence Number is a local matter, but MUST allow the BGP speaker to unambiguously identify a capability revision it initiated previously based on the Sequence Number carried in the acknowledgement from the peer.

If the Init/Ack bit is set to 1 for a capability revision in a CAPABILITY message received by a BGP speaker, then the BGP speaker SHALL treat the capability revision as an acknowledgement of the receipt of a capability revision initiated by the BGP speaker. The BGP speaker MUST ignore the Ack Request bit, and SHALL use the Sequence Number carried in the capability revision to match with the capability revision previously initiated. The BGP speaker SHALL ignore an acknowledgement for a capability revision in which an acknowledgement was not requested by the BGP speaker. If the Sequence Number carried in the capability revision does not match any of the the Sequence Numbers used in the capability revisions initiated by the BGP speaker, then the BGP speaker SHOULD send a NOTIFICATION message as specified in the Error Handling section.

If the Init/Ack bit is set to 0 for a capability revision in a CAPABILITY message received by a BGP speaker, then the BGP speaker SHOULD first validate the capability code in the message. If the capability code is not listed in the Dynamic Capability advertised by the speaker to the peer, the BGP speaker SHOULD send a NOTIFICATION message as specified in the Error Handling section. For a valid capability code, if the Ack Request bit is set to 1, the BGP speaker MUST first send a CAPABILITY message to acknowledge the receipt of the capability revision. The Init/Ack bit in the acknowledgement

MUST be set to 1, and all the other fields in the capability revision MUST be kept unchanged except that the triple <Capability Code, Capability Length, Capability Value> MAY be optionally excluded.

After receiving a capability revision initiated by a peer, the BGP speaker SHALL update the capability previously received from that peer based on the Action bit in the message, and then function in accordance with the revised capability for the peer. The procedures specified in the "Error Handling" section SHOULD be followed when an error is detected in processing the CAPABILITY message.

In order to avoid ambiguities in sending and processing UPDATE messages, certain capability revisions may require close coordination between the BGP speaker (the Initiator) that initiates the capability revisions and another BGP speaker (the Receiver) that receives the capability revisions. The mechanism of acknowledgement defined in this document SHALL be used for the revision of such a capability. For the Initiator, the capability revision SHALL take effect (for sending updates) immediately after the capability revision is sent, and the capability revision SHALL take effect (for receiving updates) immediately after an acknowledgement is received from the Receiver. For the Receiver, the capability revision SHALL take effect (for receiving updates) immediately after the capability revision is received from the Initiator, and the capability revision SHALL take effect (for sending updates) immediately after an acknowledgement is sent.

8. Error Handling

This document defines a new NOTIFICATION error code:

Error Code	Symbolic Name
7	CAPABILITY Message Error

The following error subcodes are defined as well:

Subcode	Symbolic Name
1	Unknown Sequence Number
2	Invalid Capability Length
3	Malformed Capability Value
4	Unsupported Capability Code

If a BGP speaker detects an error while processing a CAPABILITY

message, it MUST send a NOTIFICATION message with Error Code CAPABILITY Message Error. If any of the defined error subcode is applicable, the Data field of the NOTIFICATION message MUST contain the tuple for the capability revision that causes the speaker to send the message.

If the Sequence Number carried in a capability revision marked as acknowledgement does not match any of the the Sequence Numbers used in the capability revisions initiated by the BGP speaker, then the error subcode is set to Unknown Sequence Number.

If the Capability Length field in the CAPABILITY message is incorrect for a Capability Code, then the error subcode is set to Invalid Capability Length.

If the Capability Value field in the CAPABILITY message is malformed (the definition of "malformed" depends on the Capability Code), then the error subcode is set to Malformed Capability Value.

If the Capability Code in the CAPABILITY message is not any of the capability codes advertised in the Dynamic Capability by the speaker, then the error subcode is set to Unsupported Capability Code.

9. IANA Considerations

This document uses a BGP capability code to indicate that a BGP speaker supports the Dynamic Capability. The capability code needs to be assigned by IANA per [RFC 2842](#).

10. Security Considerations

This extension to BGP does not change the underlying security issues [[BGP-MD5](#)].

11. Acknowledgments

The authors would like to thank Yakov Rekhter, Ravi Chandra, Dino Farinacci, Pedro Marques, Chandrashekhara Appanna, Derek Yeung, Bruno Rijnsman and John Scudder for their review and comments.

12. References

[BGP-4] Rekhter, Y., T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [draft-ietf-idr-bgp4-24.txt](#), November 2003.

[BGP-MP] T. Bates, R. Chandra, D. Katz, and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 2858](#), June 2000.

[BGP-CAP] R. Chandra, J. Scudder, "Capabilities Advertisement with BGP-4", [RFC 2842](#), May 2000.

[BGP-MD5] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", [RFC 2385](#), August 1998.

[RFC-2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC-2026] Bradner, S., "The Internet Standards Process -- Revision 3", [RFC 2026](#), October 1996.

13. Author Information

Enke Chen
Redback Networks, Inc.
300 Holger Way
San Jose, CA 95134
e-mail: enke@redback.com

Srihari R. Sangli
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134
e-mail: rsrihari@cisco.com

14. Intellectual Property Considerations

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.

15. Full Copyright Notice

Copyright (C) The Internet Society (2004). 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 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.

