

Dynamic Capability for BGP-4
draft-ietf-idr-dynamic-cap-12.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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/1id-abstracts.html>

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

This Internet-Draft will expire on April 13, 2011.

Copyright Notice

Copyright (c) 2010 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.

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.

1. Introduction

Currently BGP capabilities [[RFC5492](#)] 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 [[RFC4760](#)]), 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.

1.1. 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 [[RFC2119](#)].

2. Dynamic Capability

The Dynamic Capability is a new BGP capability [[RFC5492](#)]. The Capability Code for this capability is specified in the "IANA Considerations" section of this document. The Capability Value field consists of a list of capability codes (one-octet for each) that specify the capabilities that MAY be revised dynamically by the remote 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.

3. Capability Message

The CAPABILITY Message is a new BGP message type with type code 6. In addition to the fixed-size BGP header [[RFC4271](#)], 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 acknowledgment 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 acknowledgment 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 [[RFC5492](#)], and it specifies a capability for which the "Action" shall be applied.

4. 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 [[RFC5492](#)] 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 [[RFC4271](#)]. 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 acknowledgment 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 acknowledgment 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 acknowledgment for a capability revision in which an acknowledgment 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 acknowledgment MUST

be set to 1, and all the other fields in the capability revision MUST be kept unchanged.

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 BGP speaker SHALL ignore such a capability revision that either results in no change to an existing capability, or removes a capability that was not advertised previously. 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 acknowledgment 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 acknowledgment 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 acknowledgment is sent.

5. 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 acknowledgment 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.

6. IANA Considerations

This document defines the CAPABILITY message type for BGP with type code 6, and a NOTIFICATION error code and subcodes for the errors in a CAPABILITY message.

This document uses a BGP capability code to indicate that a BGP speaker supports the Dynamic Capability. The capability code 67 has been assigned by IANA.

7. Security Considerations

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

8. Acknowledgments

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

9. Normative References

[RFC4271] Rekhter, Y., T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

[RFC4760] Bates, T., Chandra, R., Rekhter, Y., and D. Katz, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.

[RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), February 2009.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

10. Authors' Addresses

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

Email: enkechen@cisco.com

Srihari R. Sangli
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134

Email: rsrihari@cisco.com

