

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: 5 January 2026

E. Chen
Palo Alto Networks
S. Sangli
Juniper Networks
4 July 2025

Dynamic Capability for BGP-4
draft-ietf-idr-dynamic-cap-17

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.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 5 January 2026.

Copyright Notice

Copyright (c) 2025 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

| | |
|---|----|
| 1. Introduction | 2 |
| 1.1. Requirements Language | 3 |
| 2. Enabling a Capability on a peering session | 3 |
| 2.1. Dynamic revision of a Capability via BGP Dynamic Capability | 4 |
| 3. BGP Dynamic Capability Message | 5 |
| 4. Procedures for handling BGP Dynamic Capability Revisions | 6 |
| 4.1. Procedures for the Initiator | 7 |
| 4.2. Procedures for the Receiver | 8 |
| 5. Revising capabilities via Dynamic Capability | 9 |
| 6. Limitations of the BGP Dynamic Capability | 10 |
| 7. Error Handling | 10 |
| 8. Backward compatibility with existing deployment | 11 |
| 8.1. New Initiator and Old Receiver. | 12 |
| 8.2. Old Initiator and New Receiver. | 12 |
| 9. IANA Considerations | 12 |
| 10. Security Considerations | 12 |
| 11. Acknowledgments | 12 |
| 12. Appendix | 13 |
| 12.1. Appendix 1: Changes from Version 16 | 13 |
| 12.2. Appendix 2: Summary of Major Revisions | 13 |
| 12.3. Appendix 3: Operational states for Capability Revision | 14 |
| 12.3.1. Appendix 3.1: Initiator Capability Revision | 15 |
| 12.3.2. Appendix 3.2: Receiver Capability Revision | 15 |
| 12.4. Deployment use cases | 16 |
| 12.4.1. Adding a capability dynamically | 16 |
| 12.4.2. Deleting a capability dynamically | 16 |
| 12.4.3. Network upgrade without Non-Stop-Routing | 17 |
| 12.4.4. Network upgrade with Non-Stop-Routing | 17 |
| 13. References | 17 |
| 13.1. Normative References | 17 |
| 13.2. Informative References | 18 |
| Authors' Addresses | 19 |

1. Introduction

Currently, the BGP capabilities [RFC5492] are only advertised in the BGP OPEN message [RFC4271] during the session initialization. In order to enable or disable a capability (such as the Address Family support [RFC4760]), an established session would need to be reset, which may disrupt other services running over the session. Also, an advertised capability cannot be updated on-demand over an established session. One example of such a requirement is for adjusting the "Restart Time" in the Graceful Restart Capability [RFC4724]) when performing certain planned maintenance in a network.

Most capabilities define just one instance of the capability (also known as single-instance capabilities). Certain other capabilities have multiple instances of capability (also known as multi-instance capabilities). Route Refresh capability [RFC2918], is an example for single-instance capability and the Multiprotocol extensions for BGP-4 capability [RFC2858] is a multi-instance capability as it can list one or more individual address-family, sub-address-family capabilities.

The IANA BGP protocol registry lists the capabilities that a BGP speaker can advertise during session establishment phase. It would benefit network operations if each of these capabilities can be revised dynamically without resetting 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. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Enabling a Capability on a peering session

[RFC5492] specifies Capabilities advertisement for BGP as an optional parameter in OPEN message. By announcing the capability via OPEN message a BGP speaker conveys that it is capable of receiving and properly handling messages related to that capability. A BGP speaker may publish one or more capabilities during the session establishment. This document extends the usage of capability.

By definition, the BGP capability refers to the advertising router's behavior. In order to support a capability and exchange messages for that capability on a peering session, the capabilities have to be advertised by the peering BGP speakers. Certain capabilities require both BGP speakers to advertise them, while for certain other capabilities, advertisement from only one BGP speaker is sufficient. The list of capabilities advertised by two peers may be non-congruent.

The type of capability advertised by a BGP speaker will determine its behavior during the peering session. For example, Route Refresh Capability can be advertised by only one BGP speaker and by doing so,

it interprets and handles incoming Route-Refresh messages. A BGP speaker that supports the Multiprotocol Extensions for BGP-4 capability, may use it after determining that the peer also supports this capability. if an operator wishes to enable new functionality during the lifetime of a BGP session, and if that requires a BGP speaker to revise one or more capabilities, it can do so by resetting the session and advertise the capabilities during OPEN as per [RFC5492]

2.1. Dynamic revision of a Capability via BGP Dynamic Capability

This document proposes a new BGP capability called Dynamic Capability, defined as per 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 BGP speaker during the lifetime of a session. The list of capabilities in the value field of Dynamic Capability TLV s hereby referred as DCAP list.

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 DYNAMIC CAPABILITY message (as defined in the next Section) from the peer after the BGP session has been established. A BGP speaker may announce Dynamic Capability in the OPEN message and during the lifetime of the session, and it may revise one or more capabilities or capability-instances. The BGP speaker that revises its capability by sending the Dynamic Capability message is hereby referred to as Initiator. The remote BGP speaker that responds to the Dynamic Capability message is hereby referred to as Receiver.

Via the Dynamic Capability message, a BGP speaker (Initiator) will revise its capabilities. The receiving BGP speaker (Receiver) will make a note of the Initiator's capability revisions and sends messages to the Initiator pertaining to that capability. While many capabilities enable information exchange via existing BGP messages, some require a change in the format of the message. For example, the add-path capability [RFC7911] or 4-octet AS capability [RFC6793] change the structure of BGP update messages.

This document limits the scope of the dynamic revision of capabilities and following capability revisions are allowed.

- * Only Capabilities that do not change the format of the existing messages.

- * Only individual capability-instance capabilities under multi-instance capabilities

This document describes procedures for a 2-way handshake for capability revision. Given the underlying TCP's reliable transport, the 2-way handshake procedure is sufficient to create consistent state on both Initiator and Receiver as they implement the capability revision. The Initiator will initiate the handshaking process with a capability revision Init message. The Receiver will acknowledge the capability revision by sending a Ack message. The receipt of the ack message completes the 2-way handshaking procedure and the two speakers can then put the revised capability into effect. The capabilities that do not result in change in the format of any existing message structure can be revised dynamically via 2-way handshake. The capabilities that change the format of existing message structure can be revised during OPEN message or dynamically via [I-D.chen-idr-enhanced-dynamic-cap] which proposes additional protocol procedures.

3. BGP Dynamic Capability Message

The DYNAMIC CAPABILITY (hereby referred to as DCAP) Message is a new BGP message type with type code 6. In addition to the fixed-size BGP header [RFC4271], the DCAP message contains the following fields for revising a capability:

| |
|------------------------------|
| +-----+ |
| 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) |
| +-----+ |

Table 1

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 MAY be used by a BGP speaker to co-relate the responses to the capability revision that the speaker initiated previously for debugging purposes.

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. The Capability Length field, though, is larger than the one specified in [RFC5492].

If multiple capability instances (as described in [RFC5492]) are defined for the capability code (also known as multi-instance capability), then each capability instance MUST be revised individually, one capability instance at a time. The triple <Capability Code, Capability Length, Capability Value> in the CAPABILITY message MUST contain only one instance of the capability.

For single-instance capability the "Action" specified applies to that capability identified by the capability code. Furthermore, if the "Action" is to remove a capability, then the Capability Length field SHOULD be set to zero by the sender and the Capability Value field MUST be ignored by the receiver even when the Capability Length field has a non-zero value.

4. Procedures for handling BGP Dynamic Capability Revisions

A BGP speaker that is willing to receive the DCAP message for a capability from its peer SHOULD use the BGP Capabilities Advertisement [RFC5492] to advertise the Dynamic Capability containing the capability code. A DCAP message MAY be received only in the Established state. Receiving a DCAP 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 DCAP message from its peer.

The Initiator will need a demarcation from the Receiver acting as a confirmation so it can act on the capability revision. Similarly, the Receiver will need a demarcation to act on the revised

capability. The demarcation indicator is a message sent by the remote BGP speaker. The section below specifies the demarcation indicator for the Initiator and Receiver, and their procedures.

4.1. Procedures for the Initiator

The Initiator MUST only proceed with following steps if the Receiver has advertised Dynamic Capability indicating that it is capable of handling DCAP message. For the capability 'c' that the Initiator is going to revise, it MUST also verify if the Receiver has listed capability 'c' in the DCAP list. If the Receiver is not capable of Dynamic Capability or if 'c' is not in the DCAP list, the Initiator should log and discard the capability revision.

When the Initiator sends a DCAP 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 indicating that the capability revision has been initiated. The Ack Request bit MUST be set to 1 indicating that capability MUST be acknowledged. The assignment of the Sequence Number is a local matter, and may be used to correlate the responses from the Receiver. This can be helpful during troubleshooting any problems in capability revision. The capability that is being revised will be encoded as per IANA BGP Protocol registry capability codes. While a capability revision is in progress, the Initiator MUST NOT initiate another revision of the same capability (or the same capability instance for a multi-instance capability).

After sending the DCAP message revising a capability and before processing the acknowledgement message from the Receiver, the Initiator MUST operate as if the capability revision has not been initiated. During this phase, it must continue its usual operation of sending, receiving and processing BGP messages from the Receiver BGP speaker as specified below.

- * If the Initiator intends to add a new capability, it must not enable the capability or send messages based on the new capability revision. As there is no prior state, it MUST discard any received messages pertaining to that capability.
- * If the Initiator intends to remove an existing capability, it must not disable the capability but continue to send and process the received messages pertaining to that capability.

After receiving the DCAP message carrying capability acknowledgement with Init/Ack bit set to 1 from the Receiver, the Initiator MUST validate the DCAP message verifying the Capability code that was revised. This is the demarcation indicator for the Initiator. With this, the Initiator's capability revision finite state machine is complete and it can then function in accordance with the new capability revision as follows:

- * If the Initiator added the capability, it can now process any new messages received, based on the revised capability.
- * If the capability was withdrawn by the Initiator, it may reset the internal state. Since the prior state is cleared, it may begin to discard the new messages that may be received from the Receiver pertaining to the removed capability.

To put an upper bound on the amount of time for capability revision, an implementation MUST support a (configurable) timer CapabilityRevisionTimer that imposes this upper bound. The Initiator starts the CapabilityRevisionTimer when it starts the capability revision by sending the Init message. The timer is stopped with Initiator receives the Ack message from the Receiver. When CapabilityRevisionTimer times out and the capability revision is still in progress, the dynamic capability revision MUST be discarded. This document recommends logging this error condition for troubleshooting purpose and no further attempts for dynamic capability revision should be made without administrator intervention. This document recommends 10 minute timeout value or a similar large value to avoid premature discard of capability revision.

4.2. Procedures for the Receiver

The Receiver should expect more than one capability tuple in the DCAP message and should process each capability revision independently. In the received DCAP message, if the Init/Ack bit is set to 1, it SHOULD silently discard the capability revision. For troubleshooting purposes, the unexpected acknowledgement may be logged.

If the Init/Ack bit is set to 0, the Receiver MUST first validate the capability code. If the capability code is not listed in the Dynamic Capability (DCAP list) advertised by the Receiver itself, and the Receiver MUST send a NOTIFICATION message back to the Initiator as specified in the Error Handling section. For a valid capability code, the Receiver MUST treat it as an indication of demarcation for that capability revision.

If the Ack Request bit is set to 1, the Receiver MUST send a DCAP message to acknowledge the receipt of the capability revision. The Init/Ack MUST be set to 1, indicating the acknowledgement, and all the other fields in the DCAP message MUST be kept unchanged.

The Receiver SHALL update the capability previously received from the Initiator based on the Action bit in the message, and then function in accordance with the revised capability for the peer. The Receiver 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.

5. Revising capabilities via Dynamic Capability

There is a distinction between a BGP speaker allowing a revision of one or more capabilities and a BGP speaker revising one or more capabilities. The former allows a remote BGP speaker to revise its capabilities and the local BGP speaker supports that revision. The latter is about the local router operationalizing a capability, and putting it into effect.

A BGP speaker may choose to advertise one of more capabilities. If it has advertised Dynamic Capability (via OPEN or dynamically) it can accept Dynamic Capability message from remote BGP speaker, The value of Dynamic Capability TLV is DCAP list. By having a capability in the DCAP list, the local BGP speaker is indicating that it has the support and ability to handle the revision (add or delete) of that capability. The remote BGP speaker makes a note of the list of capabilities in the DCAP list and performs the revision during the lifetime of the peering session.

It is quite possible to list Dynamic Capability itself in DCAP list. This means that local BGP speaker can handle the revision of Dynamic Capability itself, thereby allowing add/delete capabilities from DCAP list. This document recommends that BGP speakers list the Dynamic Capability Code in Dynamic Capability. This will allow a BGP speaker to revise the list capability instances during the lifetime of the peering session by sending a DCAP message with Dynamic Capability revising DCAP list.

6. Limitations of the BGP Dynamic Capability

If the capability results in change of the format of the messages, it is important to have tighter co-ordination. For example, the procedures specified in this document does not provide demarcation enough for the Receiver to know when the Initiator will advertise the messages based on the revised capability. Hence, the capabilities that have bi-directional capability dependency requiring 3-way handshake will not function accurately. With this limitation, following capabilities can be revised using the procedures mentioned in this document.

- * Multiprotocol Extensions for BGP-4
- * Route Refresh Capability for BGP-4
- * BGP Role
- * Graceful Restart
- * Enhanced Route Refresh
- * Long-Lived Graceful Restart
- * Routing Policy Distribution
- * FQDN

The remaining capabilities may only be advertised via OPEN message during session establishment.

7. Error Handling

This document defines a new NOTIFICATION error code:

Error Code: TBD

Symbolic Name: CAPABILITY Message Error

The following error subcodes are defined:

| Subcode | Description |
|---------|--------------------------------------|
| 1 | Unknown Sequence Number (deprecated) |
| 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. On the receipt of such a NOTIFICATION message, the BGP speaker should log for troubleshooting purposes. The ongoing capability revision MUST be discarded by both Initiator and Receiver. No new capability revisions can be initiated until administrator intervention.

This document revises the usage of Sequence Number. The DCAP message fields are sufficient to correlate the message between the Initiator and the Receiver, hence the Sequence Number usage is limited to diagnostic purposes. The BGP speaker MUST not validate the Sequence Number received and MUST not send NOTIFICATION message with Unknown Sequence Number. If a NOTIFICATION message with Error code Unknown Sequence Number is received, it MUST be logged for troubleshooting purposes before silently discarding it.

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.

8. Backward compatibility with existing deployment

The new protocol procedures can work with existing implementations of Initiator and Receiver. The following section describes the different scenarios.

If a BGP speaker implements the new procedures specified in this document, its referred to as "New". If a BGP speaker implements the BGP Dynamic Capability procedures as specified in draft version 16 or prior, and does not implement the new procedures specified in this document, it is referred to as "Old". In this section, if a BGP speaker sends DCAP message with Init/Ack bit set to 1, the BGP speaker is referred as sending ack. Similarly, if a BGP speaker sends DCAP message with Ack Request bit set to 1, the BGP speaker is referred to requesting ack.

8.1. New Initiator and Old Receiver.

A New Initiator revises its capability and sends DCAP message to the Receiver. The New Initiator always requests for an ack, and since the Old Receiver honors the ack request, it sends the ack in DCAP message. This will complete the capability FSM on both Initiator and Receiver.

8.2. Old Initiator and New Receiver.

An Old Initiator revises its capability and sends the DCAP message to NEW Receiver. The Old Initiator may request for an ack. The New Receiver honors it and sends the ack in DCAP message. In case the Old Initiator does not request for Ack, the New Receiver also honors the same and not send ack DCAP message. This will also complete the capability FSM on both Initiator and Receiver. However, the Old Initiator will continue to have same ambiguity as before. This ambiguity will not exist if the Initiator implements the procedures mentioned in this document.

9. 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.

IANA is requested to allocate NOTIFICATION error code for handling errors during Dynamic Capability revision.

10. Security Considerations

The extension proposed in this document does not change the underlying security or confidentiality issues inherent in the existing BGP [RFC4271].

11. Acknowledgments

The authors would like to thank Yakov Rekhter, Ravi Chandra, Dino Farinacci, Pedro Marques, Chandrashekhar Appanna, Derek Yeung, Bruno Rijnsman, John Scudder, Jeffrey Haas, Heidi Ou, Tony Przygienda, Krishnaswamy Ananthamurthy for their review and comments.

12. Appendix

This section provides additional information useful for reviewers, and operators.

12.1. Appendix 1: Changes from Version 16

The following list highlights the updates in the current version of the document and compares with [I.D. draft-ietf-idr-dynamic-cap] version 16 or prior.

| Current version | Version-16 or prior |
|---|---|
| Supports only certain capabilities for dynamic capability revision. | Supports all capabilities for dynamic capability revision. |
| Only one capability tuple can be revised in the DCAP message. | More than one capability tuple can be revised in a DCAP message. |
| Initiator always sets Ack Request to 1, asking for acknowledgement | Initiator may or may not set Ack Request to 1, making ack optional |
| Sequence number is used for troubleshooting purposes only. NOTIFICATION message not sent. | Sequence number is used for correlation and NOTIFICATION message may be sent. |
| Clarification on when capability revision must be put into effect | |
| Clarification and procedures for sending NOTIFICATION messages | |
| Clarification and procedures for handling NOTIFICATION messages | |

12.2. Appendix 2: Summary of Major Revisions

In version 03, The Capability Length field is changed from zero octet to one octet, and the Capability Value field is specified for listing the capability codes (one-octet for each) for which the dynamic revision is supported by a BGP speaker.

In version 05, the CAPABILITY message was changed and several new fields were added, including Init/Ack, Ack Request, Reserved, and Sequence Number. In addition, the old capability code (66) was deprecated, and a new capability code (67) was allocated.

In version 06, the Capability Length field in the CAPABILITY message was increased from one octet to two octets.

In version 16, several clarifications were made about multi-instance capabilities. Also the Implementation Considerations section was added.

In version 17, further clarifications are made about multi-instance capabilities. The Error Code is changed from 7 to TBD due to a conflict.

12.3. Appendix 3: Operational states for Capability Revision

A BGP speaker MUST maintain states about whether a capability has been advertised, or received during the lifetime of the BGP session. For a multi-instance capability, the states of the capability and its revision MUST be instance specific.

The following symbols are designated for that purpose:

"L:" refers to the local speaker and "R:" refers to the remote speaker

| | |
|-------------------------|--|
| L:Cap.True | - Capability advertised |
| R:Cap.True | - Capability received |
| L:Cap.False | - Capability not advertised |
| R:Cap.False | - Capability not received |
| L:Dyn.Oper.None/Add/Del | - Following Capability revision may be triggered at Idle state Operator adds a capability OR Operator deletes a capability |
| L.Send.None | - Router does not send DCAP message |
| R.Send.None | - Router does not send DCAP message |
| L.Recv.None | - Router does not receive DCAP messages |
| R.Recv.None | - Router does not receive DCAP messages |
| L.Send.Init | - Router sends DCAP message with Init/Ack=0 |
| R.Recv.Init | - Router receives DCAP message with Init/Ack=0 |
| L.Recv.Ack | - Router receives DCAP message with Init/Ack=1 |
| R.Send.Ack | - Router sends DCAP message with Init/Ack=1 |

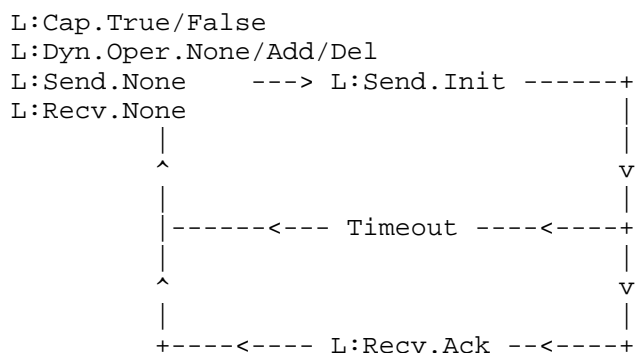
During the dynamic revision of a capability, there are separate states, "Sending State", and "Receiving State" driven by the Dynamic Capability revision.

12.3.1. Appendix 3.1: Initiator Capability Revision

States for Initiator as it revises its capability.

```
L.Dyn.Oper.None/Add/Del
L:Send.None
L:Recv.None
L:Send.Init
L:Recv.Ack
```

State transitions:



12.3.2. Appendix 3.2: Receiver Capability Revision

States for Receiver as it receives the capability revision.

```

R:Dyn.Oper.None/Add/Del
R:Recv.None
R:Send.None
R:Recv.Init
R:Send.Ack

```

State transitions:

```

R:Cap.True/False
R:Dyn.Oper.None/Add/Del
R:Send.None      ---> R:Recv.Init ---+
R:Recv.None      |
                  |
                  ^
                  |
                  +-----<----- R:Send.Ack -----+

```

12.4. Deployment use cases

The Initiator and Receiver may advertise one or more capabilities via OPEN. They must also advertise Dynamic Capability via OPEN. With this, the BGP speakers can enable revision of any capability dynamically as described in the following sections.

12.4.1. Adding a capability dynamically

During session establishment, the Receiver advertises Cap-1, Cap-2 and DynCap in OPEN message. The DynCap list contains Cap-1 and Cap-2 indicating that Receiver is capable of handling dynamic revision of capabilities Cap-1 and Cap-2 dynamically. The Initiator during session establishment advertises Cap-1 and DynCap in OPEN message. After session is established, sometime during the lifetime of the peering session, the Operator enables a functionality on the Initiator that requires Cap-2. Since Receiver allows dynamic revision of Cap-2, the Initiator sends DCAP message with Action "add" for Cap-2. The Receiver acknowledges addition of Cap-2 and after receiving the ack, the Initiator puts additions of Cap-2 into effect.

12.4.2. Deleting a capability dynamically

During session establishment, the Initiator advertises Cap-1, Cap-2 and DynCap in OPEN message. Similarly, the Receiver advertises Cap-1, Cap-2 and DynCap in OPEN message. The DynCap list contains Cap-1 indicating that Receiver is capable of handling dynamic revision of Cap-1 only. During the lifetime of the peering session, the operator removes a functionality due to which Initiator sends DCAP message with Action "delete" for Cap-1. The Receiver acknowledges deletion

of Cap-1 and after receiving the ack, the Initiator puts deletion of Cap-1 into effect.

12.4.3. Network upgrade without Non-Stop-Routing

The Initiator is upgraded and during session establishment, it advertises Cap-1 and DynCap in OPEN message. The Receiver is not upgraded and during the session establishment, it advertises Cap-1 and DynCap in OPEN message. The DynCap lists Cap-1 indicating that Receiver is capable of handling dynamic revision of Cap-1 only. As part of the upgrade, the Initiator supports new functionality and new capability Cap-2. The Operator wishes to enable new functionality during the lifetime of the peering session and as a result the Initiator wants to "add" Cap-2. It cannot send DCAP message with Action "add" for Cap-2 because the Receiver does not have support to handle Cap-2. The Initiator can add Cap-2 only when Receiver allows the dynamic revision of Cap-2.

12.4.4. Network upgrade with Non-Stop-Routing

The Initiator has advertised Cap-1 and DynCap capabilities in the OPEN message. The Receiver has advertised Cap-1 and DynCap capabilities in the OPEN message. The DynCap lists contains Cap-1 and DynCap capabilities indicating that Receiver is capable of handling dynamic revision of Cap-1 and DynCap list. The Initiator is upgraded without resetting the BGP session and the Initiator now has additional capability Cap-2 that it wants to revise. The Initiator will revise DynCap capability with action "add" and the DynCap list will contain Cap-1, Cap-2 and DynCap capabilities. The receiver acknowledges that revision of DynCap capability. When the Receiver is updated without resetting the session, it may revise the DynCap capability adding Cap-1, Cap-2 and DynCap capabilities. With this, the Receiver is indicating that it is capable of handling Cap-2 capability revision. The Initiator may send DCAP message with Action "add" for Cap-2. The Receiver may follow the similar process independently, thus allowing asynchronous network upgrade without resetting the BGP session.

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <<https://www.rfc-editor.org/info/rfc5492>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

13.2. Informative References

- [RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", RFC 4724, DOI 10.17487/RFC4724, January 2007, <<https://www.rfc-editor.org/info/rfc4724>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.
- [RFC2858] Bates, T., Rekhter, Y., Chandra, R., and D. Katz, "Multiprotocol Extensions for BGP-4", RFC 2858, DOI 10.17487/RFC2858, June 2000, <<https://www.rfc-editor.org/info/rfc2858>>.
- [RFC8654] Bush, R., Patel, K., and D. Ward, "Extended Message Support for BGP", RFC 8654, DOI 10.17487/RFC8654, October 2019, <<https://www.rfc-editor.org/info/rfc8654>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <<https://www.rfc-editor.org/info/rfc6793>>.
- [RFC2918] Chen, E., "Route Refresh Capability for BGP-4", RFC 2918, DOI 10.17487/RFC2918, September 2000, <<https://www.rfc-editor.org/info/rfc2918>>.

[I-D.chen-idr-enhanced-dynamic-cap]

Chen, E. and S. R. Sangli, "Enhanced Dynamic Capability for BGP", Work in Progress, Internet-Draft, draft-chen-idr-enhanced-dynamic-cap-00, 28 March 2025, <<https://datatracker.ietf.org/doc/html/draft-chen-idr-enhanced-dynamic-cap-00>>.

Authors' Addresses

Enke Chen
Palo Alto Networks
3000 Tannery Way
Santa Clara, CA 95054
USA
Email: enchen@paloaltonetworks.com

Srihari Sangli
Juniper Networks
Exora Business Park
Bangalore, KA 560103
India
Email: ssangli@juniper.net