

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: 27 August 2026

K. Ananthamurthy
M. Mishra
L. Krattiger
Cisco
K. Patel
Arrcus
J. Haas
Juniper Networks
D. Schatte
Charter Communications
23 February 2026

BGP Route Type Capability
draft-kriswamy-idr-route-type-capability-03

Abstract

BGP supports different route types, which define the encoding of Network Layer Reachability Information (NLRI) for some of the Address Family Identifier (AFI)/Subsequent Address Family Identifier (SAFI), like Ethernet VPN (EVPN), Multicast VPN(MVPN) and so on. BGP speaker MAY reset the BGP session if a given route type in an NLRI is not supported instead of treat-as-withdraw. This document defines Optional Capabilities pertaining to the exchange of route types that are supported for a particular AFI and SAFI. This framework facilitates the maintenance of sessions without necessitating resets or the inadvertent dropping of updates.

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.

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 27 August 2026.

Copyright Notice

Copyright (c) 2026 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
2. Requirements Language	3
3. Route Type Capability	3
4. Operation	4
5. Error Handling	5
6. Security Considerations	5
7. IANA Considerations	5
8. Normative References	5
Contributors	6
Authors' Addresses	6

1. Introduction

BGP Ethernet VPN (EVPN), Multicast VPN (MVPN), and other address families are designed to accommodate multiple route types. Each route type possesses distinct key lengths that consist of several fields, which together form a BGP Route Key. In instances where a specific route type is unsupported by a BGP speaker, there is a possibility that the session may be reset. This occurs due to the inability to ascertain the key; alternatively, the speaker might interpret the update as a withdrawal and subsequently discard it.

The sending BGP speaker remains unaware when a peer drops an Update due to an unsupported route type.

This document defines an optional capability exchange for route types linked to each AFI/SAFI and BGP speakers can signal the supported route types. Send a specific route type in the MP_REACH_NLRI attribute only if they have received confirmation from their peers that the route type is supported. This process reduces the risk of unsupported route type exchanges, improving the efficiency and reliability of routing between network peers.

2. 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.

3. Route Type Capability

The Route Type Capability is a new optional BGP capability [RFC5492] that can be used by BGP speaker to indicate the route types of a given AF/SAFI supported.

This capability is defined as follows:

- * Capability code: To be assigned by IANA
- * Capability length: Variable
- * Capability value: Consists of 0 to 63 of the tuples "AFI, SAFI, Route Type length and Route Type values for address family"

+-----+
Address Family Identifier (2 octets)
+-----+
Subsequent Address Family Identifier (1 octet)
+-----+
Route Type length (1 cotet)
+-----+
Route Types (variable length)
+-----+

- * The AFI and SAFI, taken in combination, indicate that Route Types supported for routes that are advertised with the same AFI and SAFI. Route Types may be explicitly associated with a AFI and SAFI using the encoding of [RFC4760].

- * Route Type length: This field specifies the total length in octets ranging from 1 to 32, of Route Types field.
- * Route Type: This field specifies the route types and consists of a bit-string, with each bit set to indicate that the corresponding route type that can be accepted by the BGP Speaker advertising this capability.

Example encoding for Capability Value:

```

0               1
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+
|0|1|1|1|1|1|1|1|1|1|1|1|1|1|1|
+---+---+---+---+---+---+---+---+

```

Bit 0 set to 0, while bits 1 through 15 are all set to 1, which indicates that the speaker supports route types 1 through 15. This indicates that the speaker has the capability to support route types ranging from 1 to 15 for a particular AFI and SAFI. Each of these bits corresponds to a specific route type, and setting them to 1 signifies that the speaker can successfully exchange route information of these types with its peers. Thus, the configuration effectively communicates the speaker's readiness to handle various route types within the designated AFI/SAFI.

A speaker that supports multiple "AFI, SAFI" tuples that requires route type capability exchanges includes them as multiple capabilities in the Capabilities Optional Parameter.

If route type capabilities are exchanged between BGP speakers for a specific AFI/SAFI, and the BGP peer does not support a particular route type, that route should not be advertised to the peer.

4. Operation

When the Route Type Capability is advertised, any route type bits that are not included in the encoded bit-string will be interpreted as if they were transmitted with a value of zero for the corresponding bit.

The Route Type length specifies the octets required to represent the Route Type Bits. Although the complete representation of all possible route types only demands 32 octets, it is recommended that implementations limit the number of octets transmitted to those essential for encoding the final one-bit. This practice helps optimize data transmission. Furthermore, in certain implementations, the available capacity for BGP Route Type may be limited due to the

need to convey multiple route type capabilities, potentially affecting the overall size and efficiency of the routing information exchanged. (See [I-D.ietf-idr-ext-opt-param] for discussion on a feature to address this point.)

Bit-values 0 and 255 SHOULD be set to zero as they are RESERVED.

A BGP Speaker that has received the Route Type Capability Bits MUST ensure it does not originate or propagate any BGP Route Type encoded NLRI containing a route type that is absent from the received bit-string.

A BGP Speaker that has received an AFI/SAFI without Route Type Capability MUST treat the absence as equivalent to having received the Route Type Capability Bits covered by the base or relevant specification for that address family.

The Route Type Capability MAY be dynamically updated through the use of the Dynamic Capability capability and associated mechanisms defined in [I-D.ietf-idr-dynamic-cap].

5. Error Handling

If a BGP Speaker sends BGP Route Type Capability Bits for a given AFI/SAFI and receives a BGP NLRI with an unadvertised route type, it MUST discard those routes unless specified otherwise for that address family.

6. Security Considerations

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

7. IANA Considerations

IANA is requested to assign a new BGP Capability to the Capability Codes registry from the First Come, First Served pool. The Reference for the registration is this document. The Change Controller is IETF.

8. Normative References

[I-D.ietf-idr-dynamic-cap]

Chen, E. and S. R. Sangli, "Dynamic Capability for BGP-4", Work in Progress, Internet-Draft, draft-ietf-idr-dynamic-cap-19, 15 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-dynamic-cap-19>>.

[I-D.ietf-idr-ext-opt-param]

Chen, E. and J. Scudder, "Extended Optional Parameters Length for BGP OPEN Message", Work in Progress, Internet-Draft, draft-ietf-idr-ext-opt-param-13, 22 April 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-ext-opt-param-13>>.

[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>>.

[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>>.

Contributors

In addition to the authors listed on the front page, the following coauthors have also contributed to this document:

Satya Mohanty

Authors' Addresses

Krishnaswamy Ananthamurthy
Cisco
170 W. Tasman Drive
San Jose, CA 95134
United States of America
Email: kriswamy@cisco.com

Mankamana Mishra
Cisco
170 W. Tasman Drive
San Jose, CA 95134
United States of America
Email: mankamis@cisco.com

Lukas Krattiger
Cisco
170 W. Tasman Drive
San Jose, CA 95134
United States of America
Email: lkrattig@cisco.com

Keyur Patel
Arrcus
Email: keyur@arrcus.com

Jeffrey Haas
Juniper Networks
Email: jhaas@juniper.net

Daniel Schatte
Charter Communications
Email: Daniel.Schatte@charter.com