

Internet Engineering Task Force
Internet-Draft
Updates: 6790, 7447 (if approved)
Intended status: Standards Track
Expires: 6 May 2026

B. Wen
Comcast
K. Wang
J. G. Scudder, Ed.
HPE
S. Mohanty
Zscaler
S. Krier
Cisco Systems
K. Kompella
HPE
B. Decraene, Ed.
Orange
2 November 2025

BGP Entropy Label Characteristic
draft-ietf-idr-elc-00

Abstract

The BGP Next Hop Dependent Characteristics Attribute (NHC) provides a way for a BGP speaker to advertise certain characteristics of routes. In particular, it is useful to advertise forwarding plane features.

This specification defines an NHC characteristic that can be used to advertise the ability to process the MPLS Entropy Label as an egress LSR for all NLRI advertised in the BGP UPDATE. It updates RFC 6790 and RFC 7447 concerning this BGP signaling.

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 6 May 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. Entropy Label Characteristic (ELCv3)	3
2.1. Encoding	3
2.2. Sending the ELCv3	4
2.2.1. Aggregation	5
2.3. Receiving the ELCv3	5
2.4. ELCv3 Error Handling	5
3. Legacy ELC	5
4. IANA Considerations	6
5. Security Considerations	6
6. References	6
6.1. Normative References	6
6.2. Informative References	7
Acknowledgements	8
Contributors	8
Authors' Addresses	8

1. Introduction

[I-D.ietf-idr-nhc] provides a way for a BGP speaker to advertise certain characteristics of routes. In particular, it is useful to advertise forwarding plane features.

This specification defines an NHC characteristic, called "ELCv3", to advertise the ability to process the Multiprotocol Label Switching (MPLS) Entropy Label as an egress Label Switching Router (LSR) for all NLRI advertised in the BGP UPDATE. It updates [RFC6790] and [RFC7447] with regard to this BGP signaling; this is further discussed in Section 2. ELCv3 is only relevant to NLRI of labeled address families. (The term "labeled address family" is defined in the first paragraph of Section 3.5 of [RFC9012]. In this document, we use the term "labeled NLRI" as a short form of "NLRI of a labeled address family.")

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. Entropy Label Characteristic (ELCv3)

[I-D.ietf-idr-nhc] defines the NHC as a container for characteristic TLVs. The Entropy Label Characteristic is one such characteristic.

When BGP [RFC4271] is used for distributing labeled NLRI as described in, for example, [RFC8277], the route may include the ELCv3 as part of the NHC. The inclusion of this characteristic with a route indicates that the egress of the associated Label Switched Path (LSP) can process entropy labels as an egress LSR for that route -- see Section 4.1 of [RFC6790]. Below, we refer to this for brevity as being "EL-capable."

For historical reasons, this characteristic is referred to as "ELCv3", to distinguish it from the prior Entropy Label Capability (ELC) defined in [RFC6790] and deprecated in [RFC7447], and the ELCv2 described in [I-D.scudder-bgp-entropy-label].

This section (and its subsections) replaces Section 5.2 of [RFC6790], which was previously deprecated by [RFC7447].

2.1. Encoding

The ELCv3 has characteristic code 1, characteristic length 0, and carries no value:

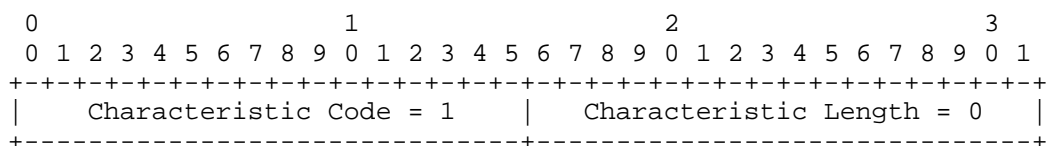


Figure 1: ELCv3 TLV Format

2.2. Sending the ELCv3

When a BGP speaker *S* has a route *R* it wishes to advertise with next hop *N* to its peer, it MAY include the ELCv3 characteristic if it knows that the egress of the associated LSP *L* is EL-capable; otherwise, it MUST NOT include the ELCv3 characteristic. Specific conditions where *S* would know that the egress is EL-capable are if *S*:

- * Is itself the egress, and knows itself to be EL-capable, or
- * Is re-advertising a BGP route it received with a valid ELCv3 characteristic, and is preserving the value of *N* as received, or
- * Is re-advertising a BGP route it received with a valid ELCv3 characteristic, and is changing the next hop that it has received to *N*, and knows that this new next hop (normally itself) is EL-capable, or
- * Is re-advertising a BGP route it received with a valid ELCv3 characteristic, and is changing the next hop that it has received to *N*, and knows (for example, through configuration) that the new next hop (normally itself) even if not EL-capable will simply swap labels without popping the BGP-advertised label stack and processing the label below, as with a transit LSR, or
- * Knows by implementation-specific means that the egress is EL-capable, or
- * Is redistributing a route learned from another protocol, and that other protocol conveyed the knowledge that the egress of *L* was EL-capable. (For example, this might be known through the Label Distribution Protocol (LDP) ELC TLV, Section 5.1 of [RFC6790].)

The ELCv3 MAY be advertised with routes that are labeled, such as those using SAFI 4 [RFC8277]. It MUST NOT be advertised with unlabeled routes.

2.2.1. Aggregation

When forming an aggregate (see Section 2.2.2 of [I-D.ietf-idr-nhc]), the aggregate route thus formed MUST NOT include the ELCv3 unless each constituent route would be eligible to include the ELCv3 according to the criteria given above.

2.3. Receiving the ELCv3

(Below, we assume that "includes the ELCv3" implies that the containing NHC has passed the checks specified in Section 2.3 of [I-D.ietf-idr-nhc]. If it had not passed, then the NHC would have been discarded and the ELCv3 would be deemed not to have been included.)

When a BGP speaker receives an unlabeled route that includes the ELCv3, it MUST discard the ELCv3.

When a BGP speaker receives a labeled route, if it includes the ELCv3, that indicates it can safely insert an entropy label into the label stack of the associated LSP. This implies that the receiving BGP speaker, if acting as ingress, MAY insert an entropy label as per Section 4.2 of [RFC6790].

2.4. ELCv3 Error Handling

The ELCv3 is considered malformed and must be disregarded if its length is other than zero.

If more than one instance of the ELCv3 is included in an NHC, instances beyond the first MUST be disregarded.

3. Legacy ELC

The ELCv3 functionality introduced in this document replaces the "BGP Entropy Label Capability Attribute" (ELC attribute) that was introduced by [RFC6790], and deprecated by [RFC7447]. The latter RFC specifies that the ELC attribute, BGP path attribute 28, "MUST be treated as any other unrecognized optional, transitive attribute". This specification revises that requirement.

As the current specification was developed, it became clear that due to incompatibilities between how the ELC attribute is processed by different fielded implementations, the most prudent handling of attribute 28 is not to propagate it as an unrecognized optional, transitive attribute, but to discard it. Therefore, this specification updates [RFC7447] by instead requiring that an implementation that receives the ELC attribute MUST discard any received ELC attribute.

4. IANA Considerations

[I-D.ietf-idr-nhc] created a new registry called "BGP Next Hop Dependent Characteristic Codes" within the Border Gateway Protocol (BGP) Parameters group and seeded it with values including:

Value	Description	Reference	Change Controller
1	ELCv3	(this doc)	IETF

Table 1

IANA is requested to update the reference to this document.

5. Security Considerations

Insertion of an ELCv3 by an attacker could cause forwarding to fail. Deletion of an ELCv3 by an attacker could cause one path in the network to be overutilized and another to be underutilized. However, we note that an attacker able to accomplish either of these (below, an "on-path attacker") could equally insert or remove any other BGP path attribute or message. The former attack described above denies service for a given route, which can be accomplished by an on-path attacker in any number of ways, even absent ELCv3. The latter attack defeats an optimization but nothing more; it seems dubious that an attacker would go to the trouble of doing so rather than launching some more damaging attack.

6. References

6.1. 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/rfc/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/rfc/rfc4271>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/rfc/rfc6790>>.
- [RFC7447] Scudder, J. and K. Kompella, "Deprecation of BGP Entropy Label Capability Attribute", RFC 7447, DOI 10.17487/RFC7447, February 2015, <<https://www.rfc-editor.org/rfc/rfc7447>>.
- [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/rfc/rfc8174>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/rfc/rfc9012>>.
- [I-D.ietf-idr-nhc]
Decraene, B., Kompella, K., Krier, S., Satya, M. R., Scudder, J., Wang, K., and B. Wen, "BGP Next Hop Dependent Characteristics Attribute", Work in Progress, Internet-Draft, draft-ietf-idr-nhc-00, 2 November 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-nhc-00>>.

6.2. Informative References

- [I-D.ietf-idr-next-hop-capability]
Decraene, B., Kompella, K., and W. Henderickx, "BGP Next-Hop dependent capabilities", Work in Progress, Internet-Draft, draft-ietf-idr-next-hop-capability-08, 8 June 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-next-hop-capability-08>>.
- [I-D.scudder-bgp-entropy-label]
Scudder, J. and K. Kompella, "BGP Entropy Label Capability, Version 2", Work in Progress, Internet-Draft, draft-scudder-bgp-entropy-label-00, 28 April 2022, <<https://datatracker.ietf.org/doc/html/draft-scudder-bgp-entropy-label-00>>.

[RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/rfc/rfc8277>>.

Acknowledgements

The authors of this specification thank Randy Bush, Mach Chen, Wes Hardaker, Jeff Haas, Susan Hares, Ketan Talaulikar, and Gyan Mishra for their review and comments.

This specification derives from two earlier documents, [I-D.ietf-idr-next-hop-capability] and [I-D.scudder-bgp-entropy-label].

[I-D.ietf-idr-next-hop-capability] included the following acknowledgements:

The Entropy Label Next-Hop Capability defined in this document is based on the ELC BGP attribute defined in section 5.2 of [RFC6790].

The authors wish to thank John Scudder for the discussions on this topic and Eric Rosen for his in-depth review of this document.

The authors wish to thank Jie Dong and Robert Raszuk for their review and comments.

[I-D.scudder-bgp-entropy-label] included the following acknowledgements:

Thanks to Swadesh Agrawal, Alia Atlas, Bruno Decraene, Martin Djernaes, John Drake, Adrian Farrell, Keyur Patel, Toby Rees, and Ravi Singh, for their discussion of this issue.

Contributors

James Uttaro
Independent Contributor
Email: juttaro@ieee.org

Wim Henderickx
Nokia
Email: wim.henderickx@nokia.com

Authors' Addresses

Bin Wen
Comcast
Email: Bin_Wen@comcast.com

Kevin Wang
HPE
Email: kfwang@juniper.net

John G. Scudder (editor)
HPE
Email: jgs@bgp.nu

Satya Mohanty
Zscaler
Email: smohanty@zscaler.com

Serge Krier
Cisco Systems
Email: sekrier@cisco.com

Kireeti Kompella
HPE
Email: kireeti@juniper.net

Bruno Decraene (editor)
Orange
Email: bruno.decraene@orange.com