

BESS Workgroup
Internet-Draft
Intended status: Standards Track
Expires: 19 November 2026

S. Sethuram
M. Mishra
Cisco Systems
S R. Mohanty
Zscaler
I. Means
P. Ramadenu
AT&T Labs, Inc.
18 May 2026

Signaling Alternative Labels for BGP Prefixes
draft-smm-bess-alt-label-01

Abstract

A single MPLS label or a label stack is advertised along with an address prefix as part of the NLRI belonging to labeled address families such as Labeled Unicast (RFC 8277), VPN (RFC 4364), etc.

This document specifies a method to advertise one or more additional, alternative labels for a set of address prefixes using the Next Hop Dependent Characteristics (NHC) attribute.

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 19 November 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. Alternative-Label Next Hop Characteristic	3
2.1. Operational behavior	4
2.1.1. Sending Alt-Label Characteristic	4
2.1.2. Receiving Alt-Label Characteristic	4
3. Extra-domain Label Type	5
3.1. Sending Extra-domain label	5
3.2. Receiving Extra-domain label	5
3.3. Application of Extra-domain labels	6
4. Error Handling	8
4.1. Extra-domain Label Type	8
5. IANA Considerations	9
6. Security Considerations	9
7. References	9
7.1. Normative References	9
7.2. Informative References	10
Appendix A. Acknowledgements	11
Authors' Addresses	11

1. Introduction

A single MPLS label or a label stack is advertised along with an NLRI belonging to labeled address families such as Labeled Unicast [RFC8277] [RFC4798], VPN [RFC4364] [RFC4659], etc. In certain topologies, it would be useful to advertise one or more additional, alternative labels associated with the same NLRI.

This document proposes a method to encode and send such alternative labels using BGP Next Hop Dependent Characteristic (NHC) attribute [I-D.ietf-idr-entropy-label].

2. Alternative-Label Next Hop Characteristic

The BGP Next Hop Dependent Characteristics attribute (NHC attribute) is an optional, transitive BGP path attribute defined in [I-D.ietf-idr-entropy-label] consisting of the Next Hop address and a list of Characteristic TLVs among other things.

A new Characteristic TLV Type called the Alternative-Label Characteristic is used to encode alternative labels for an NLRI. It is also referred to as Alt-Label Characteristic in this document.

The encoding of the Alt-Label Characteristic TLV is shown below:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Characteristic Code           | Characteristic Length           |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Label Type   | Reserved      |                             |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Label Value  |                             |                             |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Characteristic Code: Assigned by IANA to be [TBD].

Characteristic Length: Length of the Characteristic Value - set to 6.

Label Type: 1 octet representing the type of alternative label.
This field is assigned values from a new IANA registry called "BGP Alternative-Label Label Types".

Reserved: 2 octets reserved field for future use; MUST be set to 0 when sending and SHOULD be ignored when received.

Label Value: 3 octets encoding one MPLS Label [RFC3032].

This document specifies Label Type 1, called "Extra-domain Label", described in later sections. Future documents may specify newer Label Types and allocate them in the IANA registry mentioned above.

As depicted above, each Characteristic TLV supports encoding of a single label only i.e., label stack encoding is not supported.

2.1. Operational behavior

Unless specified otherwise in this document, advertisement and receive processing of the NHC attribute and the Alt-Label Characteristic TLVs follow the general behavior outlined in [I-D.ietf-idr-entropy-label].

This section specifies the general behavior for processing any Label Type. This behavior MAY be overridden for specific Label Types by documents that specify those Types.

2.1.1. Sending Alt-Label Characteristic

A router that originates a route may attach the NHC attribute containing one or more Alt-Label Characteristic TLVs. More than one such TLV may contain the same Label Type but with different label values.

2.1.2. Receiving Alt-Label Characteristic

On the receiving router:

- * any valid label in the received Alt-Label Characteristic TLV may be installed in the corresponding route forwarding entries as the outgoing label.
- * if more than one valid label is received, it is a matter of local policy to decide which of those labels are installed under what conditions.
- * if the received route is re-advertised by the router without changing the Next Hop, the received Alt-Label Characteristic TLVs MUST be re-advertised without modification.
- * if the received route is re-advertised by the router setting itself as the new Nexthop:
 - all received Alt-Label Characteristic TLVs SHOULD be removed from the NHC attribute.
 - one or more new labels MAY be allocated for the route locally, and MAY be advertised encoded in Alt-Label Characteristic TLVs within the NHC attribute.

3. Extra-domain Label Type

This document defines a new label type called "Extra-domain Label" which is assigned a value of Label Type 1. This label type indicates that any traffic destined to that label will be forwarded to one or more next hops that belong to domain(s) that are different from where the traffic ingressed. This behavior holds true whether such next hops are part of a prefix's bestpaths or backup paths or any other alternate paths.

What constitutes a "domain" is outside the scope of this document, and is defined based on local policy or configuration. One simple instance of a domain is an Autonomous System, in which case, as an example, an extra-domain label L that receives traffic from ASN-X would switch that traffic towards ASNs other than ASN-X. Other examples of a domain may be an IGP area within a single Autonomous System, or a domain spanning multiple Autonomous Systems, and so on.

3.1. Sending Extra-domain label

A router may allocate one or more extra-domain labels in addition to a "primary" label associated with the same route. The primary label is encoded as part of the NLRI (for example, [RFC8277] [RFC4364] etc) while the extra-domain labels would be encoded using the Alt-Label Characteristic TLVs.

3.2. Receiving Extra-domain label

On the receiving router:

- * the extra-domain label may be installed in the corresponding route forwarding entry as the outgoing label. Specifically this label can be installed when the route is installed as the backup path or the FRR path of a prefix. This is so that any traffic forwarded towards that label in the event of bestpaths failure will not loop back to this router.
- * if the received route is re-advertised without changing the Nexthop, the received label(s) MUST be re-advertised without modification.
- * if the received route is re-advertised by the router after setting itself as the new Nexthop:
 - the received label(s) MUST be removed from the NHC Alt-Label Characteristic TLV(s).

- one or more new extra-domain labels MAY be allocated locally and advertised by encoding them in Alt-Label Characteristic TLVs.

3.3. Application of Extra-domain labels

This section illustrates an use case where allocating an additional label for a prefix and advertising it using the Alt-Label Characteristic TLV within the NHC attribute can help break a continuous loop of advertisements.

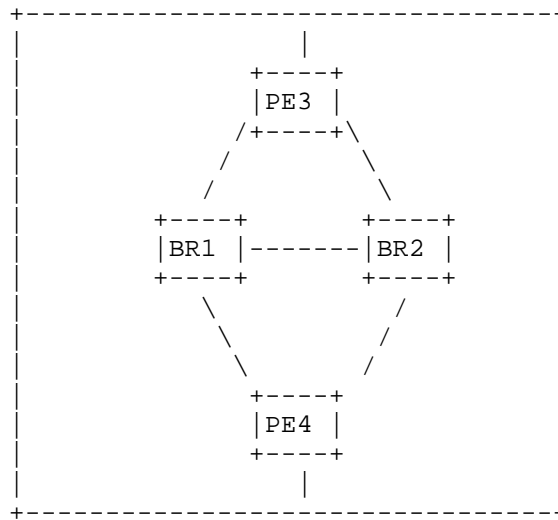


Figure 1: A pair of Border Routers backing up each other

The diagram depicts two Border Routers (BR1, BR2) peering with each other, as well as two PEs (PE3, PE4). PE3 originates some VPN routes which are advertised to BR1 and BR2, which in turn advertise them to PE4.

Step 0: PE3 originates VPN routes and advertises them to BR1 and BR2.

Step 1: Each Border Router considers its direct path to PE3 as the bestpath and selects the other Border Router as its backup path.

Step 2: Initially, BR1 receives the primary path from PE3 with label L3. Local Label L11 is allocated with context {(PE3, L3)}. This is advertised to BR2 and PE4.

Step 3: Similarly, BR2 receives the primary path from PE3 with label L3. Local Label L21 is allocated with context {(PE3, L3)}. This is advertised to BR1 and PE4.

Step 4: BR1 receives the update from BR2 and selects it to be the backup path. It now allocates a new local label L12 with context {(PE3, L3), (BR2, L21)}. This is advertised to BR2 and PE4.

Step 5: BR2 receives the new update from BR1. It now allocates a new local label L22 with context {(PE3, L3), (BR1, L12)}. This is advertised to BR1 and PE4.

Step 6: BR1 receives the new update from BR2. It now allocates a new local label L13 with context {(PE3, L3), (BR2, L22)}. This is advertised to BR2 and PE4.

Step 7: BR2 receives the new update from BR1. It now allocates a new local label L23 with context {(PE3, L3), (BR1, L13)}. This is advertised to BR2 and PE4.

The above operations continue forever resulting in a continuous loop of label allocations and advertisements.

The solution to the above problem involves each Border Router allocating an additional label whose context points to only the bestpath without including any backup path information.

Accordingly, the modified steps are described below:

Step 4: BR1 receives the update from BR2 and selects it to be the backup path. It now allocates a new local label L12 with context {(PE3, L3), (BR2, L21)}. In the route advertisement to BR2 and PE4, label L12 is encoded as part of the NLRI while label L11 is encoded as an extra-domain label in Alt-Label Characteristic TLV.

Step 5: BR2 receives the update from BR1 and selects it to be the backup path. It now allocates a new local label L22 with context {(PE3, L3), (BR2, L11)}. Note that BR2 uses the extra-domain label L11 received in the Alt-Label Characteristic TLV when allocating its new local label L22. In the route advertisement to BR1 and PE4, label L22 is encoded as part of the NLRI while label L21 is encoded as an extra-domain label in Alt-Label Characteristic TLV.

Step 6: BR1 receives the new update from BR2. It now sees the label

L21 in the Alt-Label Characteristic TLV and determines no need for a change in the already allocated label (L12) context which is {(PE3, L3), (BR2, L21)}. Hence there is no new route advertisement.

4. Error Handling

In case of any errors encountered while processing Alt-Label Characteristic TLV, the error SHOULD be logged for the attention of the operator. Unless otherwise specified, the error handling specified for NHC attribute in [I-D.ietf-idr-entropy-label] should be followed.

If the Length of the Alt-Label Characteristic TLV is not 6, then the TLV MUST be parsed according to the encoded length and MUST be ignored. Subsequent TLVs of the NHC attribute SHOULD continue to be parsed normally.

It is not an error to receive an unrecognized Label Type value in this Characteristic TLV. Such labels MUST be processed according to Section 2.1.2.

If more than one instance of a TLV with the exact same Label Type and Label value fields are received, then all but the first instance MUST be ignored.

It is not an error to receive more than one TLV instance with the same Label Type but different Label values. One or more of those labels MAY be used locally, and all such labels MUST be considered for further propagation subject to the behavior specified in Section 2.1.2.

4.1. Extra-domain Label Type

When receiving the Extra-domain Label Type, if syntactic or semantic errors are encountered while parsing the Label field, then only that TLV MUST be ignored.

It is not an error to receive more than one extra-domain label. One or more of those labels MAY be used locally, and all such labels MUST be considered for further propagation subject to the behavior specified in Section 2.1.2. The exact conditions under which a particular label may be selected for installation in local route forwarding entries is beyond the scope of this document and left to the local application.

5. IANA Considerations

This document requests IANA to allocate a new type code in the in the "BGP Next Hop Dependent Characteristic Codes" registry to represent "Alternative-Label Characteristic".

This document requests IANA to create a new registry called "BGP Alternative-Label Label Types" with the following initial allocations. The remainder of the values in the registry are to be allocated based on a First Come, First Served policy.

VALUE	DESCRIPTION	REFERENCE	CHANGE CONTROLLER
-----	-----	-----	-----
0	Reserved	(this doc)	IETF
1	Extra-domain Label	(this doc)	IETF
240-254	Private use	(this doc)	IETF
255	Reserved	(this doc)	IETF

6. Security Considerations

The security considerations of the base NHC attribute described in [I-D.ietf-idr-entropy-label] continues to apply.

The labels signaled using the new Alternative-Label Characteristic TLV is tightly associated with the Next Hop signaled in the NHC attribute. If this consistency is not maintained, intentionally or otherwise, it could lead to traffic mis-direction or blackholing. An implementaion MAY provide a local configuration option to disable installation and propagation of such received labels for specific address-families, specific routes, etc.

7. References

7.1. 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>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.

- [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>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, DOI 10.17487/RFC4659, September 2006, <<https://www.rfc-editor.org/info/rfc4659>>.
- [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>>.
- [RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, DOI 10.17487/RFC4798, February 2007, <<https://www.rfc-editor.org/info/rfc4798>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [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>>.
- [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/info/rfc8277>>.
- [I-D.ietf-idr-entropy-label]
Decraene, B., Scudder, J., Kompella, K., Satya, M. R., Wen, B., Wang, K., and S. Krier, "BGP Next Hop Dependent Characteristics Attribute", Work in Progress, Internet-Draft, draft-ietf-idr-entropy-label-18, 20 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-entropy-label-18>>.

7.2. Informative References

[RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <<https://www.rfc-editor.org/info/rfc5065>>.

Appendix A. Acknowledgements

Authors' Addresses

Shyam Sethuram
Cisco Systems
Email: shyam.ioml@gmail.com

Mankamana Mishra
Cisco Systems
Email: mankamis@cisco.com

Satya Ranjan Mohanty
Zscaler
Email: smohanty@zscaler.com

Israel Means
AT&T Labs, Inc.
Email: im8327@att.com

Praveen Ramadenu
AT&T Labs, Inc.
Email: pr9637@att.com