

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 7 June 2026

P. Mohapatra
Google LLC
R. Das, Ed.
Juniper Networks, Inc.
S. Mohanty, Ed.
Zscaler
S. Krier
Cisco Systems
R.J. Szarecki
Google LLC
A. Gattani
Arista Networks
4 December 2025

BGP Link Bandwidth Extended Community
draft-ietf-idr-link-bandwidth-22

Abstract

This document defines a BGP Extended Community, the Link Bandwidth Extended Community, which carries bandwidth information to enable weighted load-balancing in multipath scenarios. It specifies the format and processing rules for this extended community type.

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 7 June 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	3
2. Link Bandwidth Extended Community	3
3. Protocol Procedures	4
3.1. Sender (Originating Link Bandwidth Extended Community)	4
3.2. Receiver (Receiving Link Bandwidth Extended Community)	5
3.3. Re-advertisement Procedures	6
3.3.1. Re-advertisement with Next hop Change	6
3.3.2. Re-advertisement with Next Hop Unchanged	6
3.4. Link Bandwidth Extended Community Arithmetic and BGP Multipath	7
4. Error Handling	7
5. IANA Considerations	7
6. Security Considerations	8
7. Operational Considerations	8
7.1. Inconsistent Deployment	8
7.2. Bandwidth Value	9
8. Contributors	9
9. Acknowledgments	9
10. References	10
10.1. Normative References	10
10.2. Informative References	10
Appendix A. Document History	11
Authors' Addresses	11

1. Introduction

Load balancing is a critical aspect of network design, enabling efficient utilization of available bandwidth and improving overall network performance. Traditional equal load-balancing routing does not account for the varying capacities of different paths. This document suggests that the bandwidth be carried in the network using one of two new extended communities [RFC4360] - the transitive and non-transitive Link Bandwidth Extended Community. The Link Bandwidth Extended Community carries the bandwidth information of a directly connected link or multi-hop/multipath nexthop as advertised by a router. This mechanism facilitates maximizing utilization of network resources.

2. Link Bandwidth Extended Community

The Link Bandwidth Extended Community is defined as a BGP extended community that carries the bandwidth information of a router, represented by BGP Next Hop, connecting to a remote network. This community can be used to inform other routers about the available bandwidth through a given route.

The Link Bandwidth Extended Community can be either transitive or non-transitive. Therefore the value of the high-order octet of the extended Type Field can be 0x00 or 0x40, respectively. The value of the low-order octet of the extended type field for this communities is 0x04.

The Global Administrator subfield in the Value Field SHOULD be set to the Autonomous System (AS) number of the router attaching the Link Bandwidth Extended Community, but MAY contain any two-octet value. If the Autonomous System number cannot be represented in two octets, AS_TRANS [RFC6793], SHOULD be used in the Global Administrator subfield. The encoding of the full four-octet ASNs is not supported by the Link Bandwidth Extended Community. Such a capability, should the operational need for it arise, may be provided by a new BGP extension. The value in the Global Administrator subfield does not affect the use or semantics of the Link Bandwidth Extended Community. This approach maintains consistency with two-octet community registries and remains operationally familiar.

The bandwidth value is expressed as 4 octets in [IEEE.754-2019] floating point format, units being bytes (not bits!) per second. It is carried in the Local Administrator subfield of the Value Field.




Figure 1: Link Bandwidth Extended Community

Type: A 1-octet field MUST be set to 0x00 or 0x40 to indicate transitive/non-transitive.

SubType: A 1-octet field that MUST be set to 0x04 to indicate 'Link-Bandwidth'.

Global Administrator sub-field: A 2-octet field represent the Autonomous System.

Local Administrator sub-field: Bandwidth value (bytes per sec) encoded as 4 octets in IEEE 754 32-bit floating point format.

3. Protocol Procedures

The procedures cover both the transitive and non-transitive variants of the Link Bandwidth Extended Community so that implementations can handle both variants, ensuring that implementations can interoperate correctly across all deployments. Please refer to Section 5 and Appendix A for more details.

3.1. Sender (Originating Link Bandwidth Extended Community)

A BGP speaker that attaches a Link Bandwidth Extended Community SHOULD be able to advertise either a transitive or a non-transitive Link Bandwidth Extended Community. Implementations SHOULD provide configuration to set the transitivity type of the Link Bandwidth Extended Community, as well as the Global Administrator and bandwidth values in the Local Administrator sub-field, using local policy. Different implementations MAY use different default values for the transitivity type of the Link Bandwidth Extended Community. The provided configuration SHOULD allow operators to override the default transitivity value as needed. Likewise, implementation SHOULD expose their default value.

An implementation MAY advertise bandwidth value as zero. An operator may, for example, set the Link Bandwidth Extended Community to zero to indicate that the path should not attract traffic during maintenance. However, as per Section 3.2, it is up to the local policy of the receiver to decide how a link bandwidth value of zero is handled.

Generally, a single Link Bandwidth Extended Community of the transitivity type desired in a deployment is attached to a route. However during transition (refer Section 7 for details), a BGP speaker MAY attach one Link Bandwidth Extended Community per transitivity (transitive/non-transitive); the bandwidth value field in both communities SHOULD be the same.

A Link Bandwidth Extended Community MAY be attached or updated for a BGP route upon receipt during Adj-RIB-In processing. The Link Bandwidth Extended Community MAY be attached or updated for a BGP route's Adj-RIB-Out entry while being advertised to a neighboring BGP speaker.

Implementations MAY provide a configuration option to send non-transitive Link Bandwidth Extended Communities on external BGP sessions.

3.2. Receiver (Receiving Link Bandwidth Extended Community)

A BGP receiver that supports the Link Bandwidth Extended Community MUST support processing of both the transitive and non-transitive types. The receiver MUST NOT flap or treat the route as malformed based on the transitivity of the Link Bandwidth Extended Community and/or BGP session type (internal vs. external).

Implementations MAY provide configuration to accept non-transitive Link Bandwidth Extended Communities from external BGP sessions.

A BGP update with an attached Link Bandwidth Extended Community with a bandwidth value of zero is valid. When all contributing paths have a non-zero value in the Link Bandwidth Extended Community, the bandwidth values of those paths (or their ratio) can be utilized as weights to enable weighted load-balancing. Details of weighted load-balancing are outside the scope of this document. Refer to [draft-ietf-bess-ebgp-dmz] which describes some of the weighted load-balancing aspects. However, in the case where the paths have a mix of zero and non-zero values, or all zero values, the behavior is determined by local policy. For example, implementations may exclude the paths with zero value from weighted load balancing formation as long as at least one path with non-zero value exists or they may fallback to equal load-balancing. The bandwidth value, however, SHOULD NOT be used as an input to the BGP best path selection process.

Between transitive and non-transitive types of Link Bandwidth Extended Communities that have the same bandwidth value, the transitivity does not matter for purpose of computing weighted load balancing or programming to FIB (Forwarding Information Base).

3.3. Re-advertisement Procedures

This section describes the procedures to be followed when a BGP speaker receives a route with an attached Link Bandwidth Extended Community and subsequently re-advertises that route.

3.3.1. Re-advertisement with Next hop Change

When a BGP speaker re-advertises a route received with Link Bandwidth Extended Community and sets the next hop to itself or to another address, it MAY do any one of the following as its default behavior -remove the Link Bandwidth Extended Community, re-advertise it unchanged, or regenerate it with an updated value. Implementations SHOULD provide a local configuration method to alter their default behavior to the other options with per-session granularity. Likewise, implementation SHOULD expose their default value.

When regenerating Link Bandwidth Extended Community, the same procedures as outlined in Section 3.1 apply. Please also refer to Section 3.4 for use in a BGP multipath environment.

3.3.2. Re-advertisement with Next Hop Unchanged

A BGP speaker that receives a route with a Link Bandwidth Extended Community and re-advertises or reflects the same without changing its next hop, SHOULD NOT change the Link Bandwidth Extended Community in any way.

3.4. Link Bandwidth Extended Community Arithmetic and BGP Multipath

In a BGP multipath environment, the bandwidth value that is sent or re-advertised MAY be calculated based on the Link Bandwidth Extended Community associated with each constituent path contributing to multipath in the Local Routing Information Base (Local-RIB). This topic is beyond the scope of this document. Refer to [draft-ietf-bess-ebgp-dmz] which describes how this could be done in specific scenarios.

4. Error Handling

If a BGP speaker receives a route with more than one Link Bandwidth Extended Communities and uses the route to compute weighted load balancing, it SHOULD use the extended community with the lowest bandwidth value, ignoring the transitivity. Implementations MAY provide configuration to change the above preference.

A negative value in a Link Bandwidth Extended Community SHOULD NOT be attached or originated by any BGP speaker. If a BGP receiver encounters a Link Bandwidth Extended Community that contains a negative link-bandwidth value, the Link Bandwidth Extended Community SHALL be ignored.

Link Bandwidth Extended Communities with a zero value MUST NOT be considered malformed.

If any of the paths lack a valid Link Bandwidth Extended Community, equal load-balancing MUST be used instead.

5. IANA Considerations

IANA is requested to update the Transitive Two-Octet AS-Specific Extended Community Sub-Types registry (Type 0x00) and Sub-Type 0x04 [IANA-ExtComm] to:

Name

transitive Link Bandwidth Extended Community

IANA is requested to update the Non-Transitive Two-Octet AS-Specific Extended Community Sub-Types registry (Type 0x40) and Sub-Type 0x04 to:

Name

non-transitive Link Bandwidth Extended Community

Both updates are to reference this document.

6. Security Considerations

This extension to BGP has similar security implications as BGP Extended Communities [RFC4360]

The Link Bandwidth Extended Community conveys bandwidth and capacity information that may be sensitive. Exporting this community outside of an administrative domain can expose private network resource details. When propagating the routes with Link Bandwidth Extended Community towards an untrusted network or outside of an administrative domain, it is recommended operators use policy to filter out this community.

7. Operational Considerations

7.1. Inconsistent Deployment

Prior deployments of the feature specified in this document have involved implementations that only understood one of the two extended community transitivity types. As a result, such implementations would ignore the other transitivity type that they don't understand. The procedures in this document govern how multiple transitivity types for bandwidth should operate.

In circumstances where networks have deployed a mixture of implementations supporting this document's procedures for both transitivity types, and older implementations that only understand one transitivity type, inconsistent behavior could result. A prime example is when a route received by a BGP speaker contains both a transitive and a non-transitive Link Bandwidth Extended Community and that BGP speaker performs an operation that updates only one of the Link Bandwidth Extended Communities, the other community may have an inconsistent value. As a result, downstream BGP speakers that may receive such routes may perform inappropriate weighted load balancing.

To mitigate such issues, when operators are aware that older implementations are present in their networks, they may wish to take actions to address such inconsistencies. One option would be to filter either at advertisement time on the older BGP speaker the unsupported transitivity type of Link Bandwidth Extended Community - if the implementation is capable of such filtering. Alternatively, a receiving BGP speaker, knowing that the sending speaker is incapable of doing such operations, could strip the Link Bandwidth Extended Community type that is unsupported by the sender.

Ideally this operational consideration is short-lived until all the routers in the network have been upgraded to implementations that consistently support the procedures in this document.

7.2. Bandwidth Value

How the bandwidth value is computed or determined is out of scope of this document. Refer to [draft-ietf-bess-ebgp-dmz] which describes how this could be done in specific scenarios. It is recommended that implementations provide mechanisms to limit the churn caused by frequently changing bandwidth values because rapid fluctuations could impact protocol stability and network operations.

8. Contributors

Kaliraj Vairavakkalai
Juniper Networks, Inc.
1133 Innovation Way,
Sunnyvale, CA 94089
United States of America
Email: kaliraj@juniper.net

Natrajan Venkataraman
Juniper Networks, Inc.
1133 Innovation Way,
Sunnyvale, CA 94089
United States of America
Email: natv@juniper.net

Rex Fernando
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
United States of America
Email: rex@cisco.com

9. Acknowledgments

The authors would like to thank Yakov Rekhter, Srihari Sangli and Dan Tappan for proposing unequal cost load balancing as one possible application of the extended community attribute. The authors would like to thank Jeff Haas for all the discussions and providing text for operational considerations.

The authors would like to thank Bruno Decraene, Robert Raszuk, Joel Halpern, Aleksi Suhonen, Randy Bush, Stephane Litkowski, Mankamana Mishra, Moshiko Nayman, Keon Vafai, Ketan Talaulikar, Yingzhen Qu, Anoop Ghanwani, Dongjie (Jimmy) and John Scudder for their comments and contributions.

10. References

10.1. Normative References

- [IEEE.754-2019]
IEEE, "IEEE Standard for Floating-Point Arithmetic", 22 July 2019, <<https://ieeexplore.ieee.org/document/8766229>>.
- [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>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [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>>.
- [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>>.

10.2. Informative References

- [draft-ietf-bess-ebgp-dmz]
Mohanty, S., "Cumulative DMZ Link Bandwidth and load-balancing", 20 July 2025, <<https://tools.ietf.org/html/draft-ietf-bess-ebgp-dmz>>.
- [IANA-ExtComm]
IANA, "IANA Registry", <<https://www.iana.org/assignments/bgp-extended-communities/bgp-extended-communities.xhtml#trans-two-octet-as>>.

Appendix A. Document History

BGP Link Bandwidth Extended Community has evolved over several versions of the IETF draft. In the earlier versions up to draft-ietf-idr-link-bandwidth-08, only the non-transitive version of Link Bandwidth Extended Community was supported. However, starting from draft-ietf-idr-link-bandwidth-09, both transitive and non-transitive versions of Link Bandwidth Extended Community are supported.

A BGP speaker (Sender or Receiver) needs to be upgraded to support the procedures defined in this document to provide full interoperability for both transitive and non-transitive versions of Link Bandwidth Extended Community. In order to simplify implementations, it is not a goal to provide interoperability by upgrading only the RR.

Authors' Addresses

Pradosh Mohapatra
Google LLC
Email: pradosh@google.com

Reshma Das (editor)
Juniper Networks, Inc.
1133 Innovation Way,
Sunnyvale, CA 94089
United States of America
Email: dreshma@juniper.net

Satya Mohanty (editor)
Zscaler
120 Holger Way,
San Jose, CA 95134
United States of America
Email: smohanty@zscaler.com

Serge Krier
Cisco Systems
Pegasus Parc, De Kleetlaan 6a
Belgium
Email: sekrier@cisco.com

Rafal Jan Szarecki
Google LLC
1160 N Mathilda Ave,
Sunnyvale, CA 94089
United States of America
Email: rszarecki@gmail.com

Akshay Gattani
Arista Networks
5453 Great America Parkway
Santa Clara, CA 95054
United States of America
Email: akshay@arista.com