

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: 28 August 2025

J. Uttaro
Individual Contributor
A. Retana
Futurewei Technologies, Inc.
P. Mohapatra
Google
K. Patel
Arrcus, Inc.
B. Wen
Comcast
24 February 2025

One Administrative Domain using BGP
draft-uttaro-idr-bgp-oad-06

Abstract

This document defines a new External BGP (EBGP) peering type known as EBGPOAD, which is used between two EBGPeers that belong to One Administrative Domain (OAD).

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 28 August 2025.

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
1.1. Requirements Language	3
2. Discussion	3
3. Operation	4
3.1. ORIGIN	5
3.2. AS_PATH	5
3.3. NEXT_HOP	5
3.4. MULTI_EXIT_DISC	5
3.5. LOCAL_PREF	6
3.6. ATOMIC_AGGREGATE	6
3.7. AGGREGATOR	6
3.8. COMMUNITIES	6
3.9. ORIGINATOR_ID	6
3.10. CLUSTER_LIST	7
3.11. MP_REACH_NLRI	7
3.12. MP_UNREACH_NLRI	7
3.13. EXTENDED COMMUNITIES	7
3.14. AS4_PATH	7
3.15. AS4_AGGREGATOR	8
3.16. PMSI_TUNNEL	8
3.17. Tunnel Encapsulation	8
3.18. Traffic Engineering	8
3.19. IPv6 Address Specific Extended Community	8
3.20. AIGP	8
3.21. PE Distinguisher Labels	9
3.22. BGP-LS Attribute	9
3.23. LARGE_COMMUNITY	9
3.24. BGPsec_PATH	9
3.25. BGP Community Container	9
3.26. Only to Customer	10
3.27. D-PATH	10
3.28. SFP	10
3.29. BFD Discriminator	10
3.30. BGP Router Capabilities	10
3.31. BGP Prefix-SID	10
3.32. ATTR_SET	11
3.33. Summary Table	11
4. Changes to the Decision Process	13
5. Deployment and Operational Considerations	13
6. IANA Considerations	14
7. Implementation Report	16

8. Security Considerations	16
9. References	17
9.1. Normative References	17
9.2. Informative References	20
Acknowledgements	21
Contributors	21
Authors' Addresses	21

1. Introduction

At each EBGp boundary, BGP path attributes are modified as per [RFC4271], which includes stripping any attributes not allowed over an EBGp session. An example is the LOCAL_PREF attribute.

Some networks span more than one autonomous system and require more flexibility in the propagation of path attributes. It is worth noting that these multi-AS networks have a common or single administrative entity. These networks are said to belong to One Administrative Domain (OAD). It is desirable to have the ability to carry any attribute across EBGp peerings when the peers belong to an OAD.

This document defines a new EBGp peering type known as EBGp-OAD, which is used between two EBGp peers that belong to an OAD. This document also defines rules for route announcement and processing for EBGp-OAD peers.

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. Discussion

Networks have traditionally been demarcated by an autonomous system/ BGP border, which correlates to an administrative boundary. This paradigm no longer serves the needs of network designers or customers due to the decoupling of the Interior Gateway Protocol (IGP) from BGP, BGP-free core in the underlay (e.g., using BGP labeled unicast [RFC8277]), the use of BGP to facilitate multiple service overlays (e.g., L2VPN, L3VPN, etc.) spanning multiple regions and AS domains, and the instantiation of customer sites on multiple content service providers (CSPs).

For example, sites in a BGP/MPLS VPN [RFC4364] may be distributed across different AS domains. In some cases, the administrator of the VPN may prefer that some attributes are propagated to all their sites to influence the BGP decision process.

3. Operation

[RFC4271] defines two types of BGP peerings used during a BGP protocol session. As part of the extensions defined in this document, EBGp peering is divided into two types:

1. EBGp as defined in [RFC4271].
2. EBGp-OAD as defined below.

The EBGp-OAD session is a BGP connection between peers in different Autonomous Systems that belong to an OAD. By default, the EBGp-OAD speakers follow the EBGp route advertisement, route processing, path attribute announcement, and processing rules as defined in [RFC4271].

EBGp-OAD peers handle receiving optional transitive attributes as specified in [RFC4271]. Unrecognized non-transitive optional attributes MUST be quietly ignored and not passed along to other BGP peers.

Unless explicitly specified, EBGp-OAD speakers are allowed to announce and receive any attribute over an EBGp-OAD session. Receiving any attribute over an EBGp-OAD session MUST NOT result in an error. For example, the ORIGINATOR_ID (Section 3.9) and the CLUSTER_LIST (Section 3.10) are not allowed over EBGp-OAD sessions.

EBGp-OAD sessions MUST comply with the behavior specified in [RFC8212]. Furthermore, the propagation of attributes not allowed over EBGp sessions (see Table 1) MUST be explicitly allowed by an Export Policy, and their reception SHOULD be explicitly allowed by an Import Policy.

An EBGp-OAD speaker MUST support four-octet AS numbers and advertise the "support for four-octet AS number capability" [RFC6793].

This section addresses all path attributes defined at the time of this writing that are not marked as "deprecated" in the "BGP Path Attributes" registry [IANA-BGP-ATTRS]. The following subsections specify the behavior for each path attribute as related to an EBGp-OAD session. Table 1 summarizes the behavior for all session types.

Documents specifying new path attributes MUST indicate whether they are allowed for each session type: IBGP, EBGp, and EBGp-OAD.

3.1. ORIGIN

The ORIGIN attribute is a well-known mandatory BGP path attribute [RFC4271] that MUST be present in UPDATE messages sent over EBGP-OAD sessions. Its origination and value MUST comply with the specification in [RFC4271].

3.2. AS_PATH

The AS_PATH attribute is a well-known mandatory BGP path attribute [RFC4271]. It SHOULD be present in UPDATE messages sent over EBGP-OAD sessions unless it has been replaced by the BGPsec_PATH attribute [RFC8205]. The origination and modification of the AS_PATH attribute MUST comply with the EBGP-related specification in [RFC4271].

3.3. NEXT_HOP

The NEXT_HOP attribute is a well-known mandatory BGP path attribute [RFC4271] that SHOULD be present in UPDATE messages sent over EBGP-OAD sessions [RFC4760]. The origination and modification of the NEXT_HOP attribute MUST comply with the EBGP-related specification in [RFC4271].

It is reasonable for members of an OAD to share a common reachability domain. In such a case, the NEXT_HOP attribute MAY be left unchanged.

3.4. MULTI_EXIT_DISC

The MULTI_EXIT_DISC attribute is an optional non-transitive BGP path attribute [RFC4271] that MAY be present in UPDATE messages sent over EBGP-OAD sessions, even if it has been received from a neighboring AS. Otherwise, the use of the MULTI_EXIT_DISC attribute MUST comply with the specification in [RFC4271].

The determination of the neighboring AS for the purpose of BGP Route Selection [RFC4271] MAY ignore the ASNs of other members of the OAD. If so, all the members of the OAD SHOULD be configured to use the same criteria. Failure to do so may result in inconsistent forwarding between members of the OAD. Care should also be taken to avoid the creation of persistent route oscillations, similar to the Type II Churn described in [RFC3345]. [RFC7964] provides solutions and recommendations to address this issue.

3.5. LOCAL_PREF

The LOCAL_PREF attribute is a well-known BGP path attribute [RFC4271] that MAY be present in UPDATE messages sent over EBGp-OAD sessions. The use of the LOCAL_PREF attribute MUST comply with the specification in [RFC4271].

3.6. ATOMIC_AGGREGATE

The ATOMIC_AGGREGATE attribute is a well-known discretionary BGP path attribute [RFC4271] that MAY be present in UPDATE messages sent over EBGp-OAD sessions. The use of the ATOMIC_AGGREGATE attribute MUST comply with the specification in [RFC4271].

3.7. AGGREGATOR

The AGGREGATOR attribute is an optional transitive BGP path attribute [RFC4271] that MAY be present in UPDATE messages sent over EBGp-OAD sessions. The use of the AGGREGATOR attribute MUST comply with the specification in [RFC4271].

3.8. COMMUNITIES

The COMMUNITIES attribute is an optional transitive BGP path attribute [RFC1997] that MAY be present in UPDATE messages sent over EBGp-OAD sessions. The advertisement semantics MUST comply with the specification in [RFC1997].

Routes with a COMMUNITIES attribute containing the well-known NO_EXPORT community [RFC1997] SHOULD NOT be advertised across an EBGp-OAD session unless allowed by explicit policy configuration. If allowed, all the members of the OAD SHOULD be configured to use the same criteria. Failure to do so may result in inconsistent forwarding between members of the OAD.

Routes with a COMMUNITIES attribute containing the well-known NO_EXPORT_SUBCONFED community [RFC1997] MUST NOT be advertised across an EBGp-OAD session.

3.9. ORIGINATOR_ID

The ORIGINATOR_ID attribute is an optional non-transitive BGP path attribute [RFC4456] that MUST NOT be advertised over an EBGp-OAD session. If received from an EBGp-OAD neighbor, it SHALL be discarded using the "attribute discard" approach [RFC7606]. An implementation MAY log an error message for further analysis.

3.10. CLUSTER_LIST

The CLUSTER_LIST attribute is an optional non-transitive BGP path attribute [RFC4456] that MUST NOT be advertised over an EBGp-OAD session. If received from an EBGp-OAD neighbor, it SHALL be discarded using the "attribute discard" approach [RFC7606]. An implementation MAY log an error message for further analysis.

3.11. MP_REACH_NLRI

The MP_REACH_NLRI attribute is an optional non-transitive BGP path attribute [RFC4760] that MAY be advertised over an EBGp-OAD session. The use of the MP_REACH_NLRI attribute MUST comply with the EBGp-related specification in [RFC4760].

It is reasonable for members of an OAD to share a common reachability domain. In such a case, the Next Hop in the MP_REACH_NLRI attribute MAY be left unchanged.

3.12. MP_UNREACH_NLRI

The MP_UNREACH_NLRI attribute is an optional non-transitive BGP path attribute [RFC4760] that MAY be advertised over an EBGp-OAD session. The use of the MP_UNREACH_NLRI attribute MUST comply with the specification in [RFC4760].

3.13. EXTENDED COMMUNITIES

The EXTENDED COMMUNITIES attribute is a transitive optional BGP path attribute [RFC4360] that MAY be advertised over an EBGp-OAD session.

Extended communities which are non-transitive across an AS boundary MAY be advertised over an EBGp-OAD session if allowed by explicit policy configuration. If allowed, all the members of the OAD SHOULD be configured to use the same criteria. For example, the Origin Validation State Extended Community, defined as non-transitive in [RFC8097], can be advertised to peers in the same OAD.

3.14. AS4_PATH

The AS4_PATH attribute is an optional transitive BGP path attribute [RFC6793] that MAY be advertised over an EBGp-OAD session. The use of the AS4_PATH attribute MUST comply with the specification in [RFC6793].

3.15. AS4_AGGREGATOR

The AS4_AGGREGATOR attribute is an optional transitive BGP path attribute [RFC6793] that MAY be advertised over an EBGp-OAD session. The use of the AS4_AGGREGATOR attribute MUST comply with the specification in [RFC6793].

3.16. PMSI_TUNNEL

The PMSI_TUNNEL attribute is an optional transitive BGP path attribute [RFC6514] that MAY be advertised over an EBGp-OAD session. The use of the PMSI_TUNNEL attribute MUST comply with the EBGp-related specification in [RFC6514].

3.17. Tunnel Encapsulation

The Tunnel Encapsulation attribute is an optional transitive BGP path attribute [RFC9012] that MAY be advertised over an EBGp-OAD session. The use of the Tunnel Encapsulation attribute MUST comply with the EBGp-related specification in [RFC9012].

3.18. Traffic Engineering

The Traffic Engineering attribute is an optional non-transitive BGP path attribute [RFC5543] that MAY be advertised over an EBGp-OAD session. The use of the Traffic Engineering attribute MUST comply with the specification in [RFC5543].

3.19. IPv6 Address Specific Extended Community

The IPv6 Address Specific Extended Community attribute is an optional transitive BGP path attribute [RFC5701] that MAY be advertised over an EBGp-OAD session.

Extended communities which are non-transitive across Autonomous Systems MAY be advertised over an EBGp-OAD session if allowed by explicit policy configuration. If allowed, all the members of the OAD SHOULD be configured to use the same criteria.

3.20. AIGP

The AIGP attribute is an optional non-transitive BGP path attribute [RFC7311] that MAY be advertised over an EBGp-OAD session. The default value of AIGP_SESSION [RFC7311] MUST be "disabled" and it MAY be "enabled" by explicit policy configuration. The use of the AIGP attribute MUST comply with the specification in [RFC7311].

3.21. PE Distinguisher Labels

The PE Distinguisher Labels attribute is an optional transitive BGP path attribute [RFC6514] that MAY be advertised over an EBGp-OAD session. The use of the PE Distinguisher Labels attribute MUST comply with the specification in [RFC6513] and [RFC6514].

3.22. BGP-LS Attribute

The BGP Link-State (BGP-LS) attribute is an optional non-transitive BGP path attribute [RFC9552] that MAY be advertised over an EBGp-OAD session. The use of the BGP-LS Attribute MUST comply with the specification in [RFC9552].

3.23. LARGE_COMMUNITY

The LARGE_COMMUNITY attribute is an optional transitive BGP path attribute [RFC8092] that MAY be advertised over an EBGp-OAD session. The use of the LARGE_COMMUNITY attribute MUST comply with the specification in [RFC8092].

3.24. BGPsec_PATH

The BGPsec_PATH attribute is an optional non-transitive BGP path attribute [RFC8205] that MAY be advertised over an EBGp-OAD session. The use of the BGPsec_PATH attribute MUST comply with the specification in [RFC8205].

3.25. BGP Community Container

The BGP Community Container attribute is an optional transitive BGP path attribute [WIDE] that MAY be advertised over an EBGp-OAD session.

In particular, communities with the T bit [WIDE] not set MAY be advertised over an EBGp-OAD session if allowed by explicit policy configuration. Communities with the T bit set MUST be advertised over an EBGp-OAD session. Communities with the C bit [WIDE] not set MUST NOT be advertised over an EBGp-OAD session. Communities with the C bit set MAY be advertised over an EBGp-OAD session if allowed by explicit policy configuration. In all cases, all the members of the OAD SHOULD be configured to use the same criteria.

3.26. Only to Customer

The Only to Customer (OTC) attribute is an optional transitive BGP path attribute [RFC9234] that MAY be advertised over an EBGp-OAD session. However, the BGP Role negotiation and OTC Attribute-based procedures specified in [RFC9234] are NOT RECOMMENDED to be used between peers using an EBGp-OAD session. If received, the OTC attribute MUST be preserved unchanged. The use and negotiation of BGP Roles between EBGp-OAD peers is outside the scope of this document.

3.27. D-PATH

The Domain Path (D-PATH) attribute is an optional transitive BGP path attribute [IPVPN] that MAY be advertised over an EBGp-OAD session. The use of the D-PATH attribute MUST comply with the specification in [IPVPN].

3.28. SFP

The Service Function Path (SFP) attribute is an optional transitive BGP path attribute [RFC9015] that MAY be advertised over an EBGp-OAD session. The use of the SFP attribute MUST comply with the specification in [RFC9015].

3.29. BFD Discriminator

The BFD Discriminator attribute is an optional transitive BGP path attribute [RFC9026] that MAY be advertised over an EBGp-OAD session. The use of the BFD Discriminator attribute MUST comply with the specification in [RFC9026].

3.30. BGP Router Capabilities

The BGP Router Capabilities attribute (RCA) is an optional transitive BGP path attribute [ENTROPY] that MAY be advertised over an EBGp-OAD session. The use of the RCA attribute MUST comply with the specification in [ENTROPY].

3.31. BGP Prefix-SID

The BGP Prefix-SID attribute is an optional transitive BGP path attribute [RFC8669] that MAY be advertised over an EBGp-OAD session. The use of the BGP Prefix-SID attribute MUST comply with the specification in [RFC8669].

3.32. ATTR_SET

The ATTR_SET attribute is an optional transitive BGP path attribute [RFC6368] that MAY be advertised over an EBGp-OAD session. The use of the ATTR_SET attribute MUST comply with the specification in [RFC6368].

3.33. Summary Table

Path Attribute	EBGP	IBGP	EBGP-OAD	Reference
ORIGIN	Mandatory	Mandatory	Mandatory	Section 3.1
AS_PATH	Optional	Optional	Optional	Section 3.2
NEXT_HOP	Optional	Optional	Optional	Section 3.3
MULTI_EXIT_DISC	Optional	Optional	Optional	Section 3.4
LOCAL_PREF	Not allowed	Mandatory	Optional	Section 3.5
ATOMIC_AGGREGATE	Optional	Optional	Optional	Section 3.6
AGGREGATOR	Optional	Optional	Optional	Section 3.7
COMMUNITIES	Optional	Optional	Optional	Section 3.8
ORIGINATOR_ID	Not Allowed	Optional	Not allowed	Section 3.9
CLUSTER_LIST	Not Allowed	Optional	Not allowed	Section 3.10
MP_REACH_NLRI	Optional	Optional	Optional	Section 3.11
MP_UNREACH_NLRI	Optional	Optional	Optional	Section 3.12

EXTENDED COMMUNITIES	Optional	Optional	Optional	Section 3.13
AS4_PATH	Optional	Optional	Optional	Section 3.14
AS4_AGGREGATOR	Optional	Optional	Optional	Section 3.15
PMSI_TUNNEL	Optional	Optional	Optional	Section 3.16
Tunnel Encapsulation	Optional	Optional	Optional	Section 3.17
Traffic Engineering	Not Allowed	Optional	Optional	Section 3.18
IPv6 Address Specific Extended Community	Optional	Optional	Optional	Section 3.19
AIGP	Optional	Optional	Optional	Section 3.20
PE Distinguisher Labels	Optional	Optional	Optional	Section 3.21
BGP-LS Attribute	Not Allowed	Optional	Optional	Section 3.22
LARGE_COMMUNITY	Optional	Optional	Optional	Section 3.23
BGPsec_PATH	Optional	Optional	Optional	Section 3.24
BGP Community Container	Optional	Optional	Optional	Section 3.25
Only to Customer	Optional	Optional	Optional	Section 3.26
D-PATH	Optional	Optional	Optional	Section 3.27
SFP	Optional	Optional	Optional	Section

				3.28	
BFD Discriminator	Optional	Optional	Optional	Section 3.29	
BGP Router Capabilities	Optional	Optional	Optional	Section 3.30	
BGP Prefix-SID	Optional	Optional	Optional	Section 3.31	
ATTR_SET	Optional	Optional	Optional	Section 3.32	

Table 1: Path Attribute Behavior

4. Changes to the Decision Process

Section 9 of [RFC4271] describes the BGP Decision Process to select routes for local forwarding and subsequent advertisement. Section 9.1.2.2 of [RFC4271] describes tie breaking procedures in cases where a BGP speaker has several routes to the same destination. This document modifies step d) as follows:

- d) If at least one of the candidate routes was received via EBGp, remove from consideration all routes that were received via EBGp-OAD and IBGP. If at least one of the candidate routes was received via EBGp-OAD, remove from consideration all routes that were received via IBGP.

The algorithm defined in [RFC5004] to avoid unnecessary path transitions between external paths MUST be used when the routes considered were received via EBGp-OAD.

5. Deployment and Operational Considerations

For the Import and Export Policies to behave as expected, both EBGp-OADGP speakers must be configured with the same session type. If only one BGP speaker is configured that way, and the other uses an EBGp session, the result is that some path attributes may be ignored and others will be discarded.

The default BGP peering type for a session that is across autonomous systems SHOULD be EBGp. A BGP implementation SHOULD provide a configuration-time option to enable the EBGp-OAD session type. The session type may be changed once the BGP connection has been established.

If multiple peerings exist between autonomous systems that belong to an OAD, all SHOULD be configured consistently. Improper configuration may result in inconsistent or unexpected forwarding. The inconsistent use of EBGP-OAD sessions is out of scope of this document.

BGP Confederations [RFC5065] provide similar behavior, on a session by session basis, as what is specified in this document. The use of confederations with an EBGP-OAD peering is out of scope of this document.

6. IANA Considerations

IANA is requested to update the BGP Path Attributes registry as shown in Table 2. Also, IANA is requested to add [this document] as a reference in the registry.

Value	Code	EBGP	IBGP	EBGP-OAD	Reference
1	ORIGIN	Mandatory	Mandatory	Mandatory	[RFC4271]
2	AS_PATH	Optional	Optional	Optional	[RFC4271] [RFC8205]
3	NEXT_HOP	Optional	Optional	Optional	[RFC4271] [RFC4760]
4	MULTI_EXIT_DISC	Optional	Optional	Optional	[RFC4271]
5	LOCAL_PREF	Not allowed	Mandatory	Optional	[RFC4271]
6	ATOMIC_AGGREGATE	Optional	Optional	Optional	[RFC4271]
7	AGGREGATOR	Optional	Optional	Optional	[RFC4271]
8	COMMUNITIES	Optional	Optional	Optional	[RFC1997]
9	ORIGINATOR_ID	Not Allowed	Optional	Not allowed	[RFC4456]
10	CLUSTER_LIST	Not Allowed	Optional	Not allowed	[RFC4456]
14	MP_REACH_NLRI	Optional	Optional	Optional	[RFC4760]
15	MP_UNREACH_NLRI	Optional	Optional	Optional	[RFC4760]

16	EXTENDED COMMUNITIES	Optional	Optional	Optional	[RFC4360]
17	AS4_PATH	Optional	Optional	Optional	[RFC6793]
18	AS4_AGGREGATOR	Optional	Optional	Optional	[RFC6793]
22	PMSI_TUNNEL	Optional	Optional	Optional	[RFC6514]
23	Tunnel Encapsulation	Optional	Optional	Optional	[RFC9012]
24	Traffic Engineering	Not Allowed	Optional	Optional	[RFC5543]
25	IPv6 Address Specific Extended Community	Optional	Optional	Optional	[RFC5701]
26	AIGP	Optional	Optional	Optional	[RFC7311]
27	PE Distinguisher Labels	Optional	Optional	Optional	[RFC6514]
29	BGP-LS Attribute	Not Allowed	Optional	Optional	[RFC9552]
32	LARGE_COMMUNITY	Optional	Optional	Optional	[RFC8092]
33	BGPsec_PATH	Optional	Optional	Optional	[RFC8205]
34	BGP Community Container	Optional	Optional	Optional	[WIDE]
35	Only to Customer	Optional	Optional	Optional	[RFC9234]
36	D-PATH	Optional	Optional	Optional	[IPVPN]
37	SFP	Optional	Optional	Optional	[RFC9015]
38	BFD Discriminator	Optional	Optional	Optional	[RFC9026]
39	BGP Router Capabilities	Optional	Optional	Optional	[ENTROPY]

40	BGP Prefix-SID	Optional	Optional	Optional	[RFC8669]	
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+
128	ATTR_SET	Optional	Optional	Optional	[RFC6368]	
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+

Table 2: BGP Path Attributes

Table 2 only includes the path attributes referenced in this document. Any Reserved, Deprecated, or Unassigned values should contain empty IBGP, EBG, and EBG-OAD columns.

7. Implementation Report

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC7942]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC7942], "This will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature".

FRRouting (<https://github.com/FRRouting/frr/pull/14537/commits/5e81120961d502a25da5a590011ed6dda86024b7>) (10.0 and above) implementation. The official documentation is available here (<https://docs.frrouting.org/en/stable-10.2/bgp.html#clcmd-neighbor-PEER-oad>).

8. Security Considerations

EBGP-OAD peering does not change the underlying security issues inherent in the existing BGP protocol, such as those described in [RFC4271] and [RFC4272]. Any security considerations related to existing path attributes apply to EBG-OAD sessions.

All BGP attributes may now be propagated to another autonomous system. However, it is expected that the new session type will only be enabled when peering with a router that also belongs to the OAD. If misconfigured, the impact is minimal due to the fact that both

[RFC4271] and [RFC7606] define mechanisms to deal with unexpected path attributes. Also, the use of the Import and Export Policy mechanisms specified in [RFC8212] are REQUIRED.

9. References

9.1. Normative References

- [ENTROPY] 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-16, 26 September 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-entropy-label-16>>.
- [WIDE] Raszuk, R., Haas, J., Lange, A., Decraene, B., Amante, S., and P. Jakma, "BGP Community Container Attribute", Work in Progress, Internet-Draft, draft-ietf-idr-wide-bgp-communities-11, 9 March 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-wide-bgp-communities-11>>.
- [IPVPN] Rabadan, J., Sajassi, A., Rosen, E. C., Drake, J., Lin, W., Uttaro, J., and A. Simpson, "EVPN Interworking with IPVPN", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-ipvpn-interworking-13, 4 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-ipvpn-interworking-13>>.
- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<https://www.rfc-editor.org/info/rfc1997>>.
- [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>>.
- [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>>.

- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<https://www.rfc-editor.org/info/rfc4456>>.
- [RFC5004] Chen, E. and S. Sangli, "Avoid BGP Best Path Transitions from One External to Another", RFC 5004, DOI 10.17487/RFC5004, September 2007, <<https://www.rfc-editor.org/info/rfc5004>>.
- [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>>.
- [RFC5543] Ould-Brahim, H., Fedyk, D., and Y. Rekhter, "BGP Traffic Engineering Attribute", RFC 5543, DOI 10.17487/RFC5543, May 2009, <<https://www.rfc-editor.org/info/rfc5543>>.
- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", RFC 5701, DOI 10.17487/RFC5701, November 2009, <<https://www.rfc-editor.org/info/rfc5701>>.
- [RFC6368] Marques, P., Raszuk, R., Patel, K., Kumaki, K., and T. Yamagata, "Internal BGP as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 6368, DOI 10.17487/RFC6368, September 2011, <<https://www.rfc-editor.org/info/rfc6368>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.
- [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>>.
- [RFC7311] Mohapatra, P., Fernando, R., Rosen, E., and J. Uttaro, "The Accumulated IGP Metric Attribute for BGP", RFC 7311, DOI 10.17487/RFC7311, August 2014, <<https://www.rfc-editor.org/info/rfc7311>>.

- [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>>.
- [RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/info/rfc7942>>.
- [RFC8092] Heitz, J., Ed., Snijders, J., Ed., Patel, K., Bagdonas, I., and N. Hilliard, "BGP Large Communities Attribute", RFC 8092, DOI 10.17487/RFC8092, February 2017, <<https://www.rfc-editor.org/info/rfc8092>>.
- [RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <<https://www.rfc-editor.org/info/rfc8205>>.
- [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>>.
- [RFC8212] Mauch, J., Snijders, J., and G. Hankins, "Default External BGP (EBGP) Route Propagation Behavior without Policies", RFC 8212, DOI 10.17487/RFC8212, July 2017, <<https://www.rfc-editor.org/info/rfc8212>>.
- [RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.
- [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/info/rfc9012>>.
- [RFC9015] Farrel, A., Drake, J., Rosen, E., Uttaro, J., and L. Jalil, "BGP Control Plane for the Network Service Header in Service Function Chaining", RFC 9015, DOI 10.17487/RFC9015, June 2021, <<https://www.rfc-editor.org/info/rfc9015>>.

- [RFC9026] Morin, T., Ed., Kebler, R., Ed., and G. Mirsky, Ed., "Multicast VPN Fast Upstream Failover", RFC 9026, DOI 10.17487/RFC9026, April 2021, <<https://www.rfc-editor.org/info/rfc9026>>.
- [RFC9234] Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages", RFC 9234, DOI 10.17487/RFC9234, May 2022, <<https://www.rfc-editor.org/info/rfc9234>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

9.2. Informative References

- [RFC3345] McPherson, D., Gill, V., Walton, D., and A. Retana, "Border Gateway Protocol (BGP) Persistent Route Oscillation Condition", RFC 3345, DOI 10.17487/RFC3345, August 2002, <<https://www.rfc-editor.org/info/rfc3345>>.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.
- [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>>.
- [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>>.
- [RFC7964] Walton, D., Retana, A., Chen, E., and J. Scudder, "Solutions for BGP Persistent Route Oscillation", RFC 7964, DOI 10.17487/RFC7964, September 2016, <<https://www.rfc-editor.org/info/rfc7964>>.
- [RFC8097] Mohapatra, P., Patel, K., Scudder, J., Ward, D., and R. Bush, "BGP Prefix Origin Validation State Extended Community", RFC 8097, DOI 10.17487/RFC8097, March 2017, <<https://www.rfc-editor.org/info/rfc8097>>.

[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>>.

[IANA-BGP-ATTRS]
IANA, "BGP Path Attributes",
<<http://www.iana.org/assignments/bgp-parameters>>.

Acknowledgements

The authors would like to thank everyone who has commented on this work, including (in alphabetical order) Donatas Abraitis, Randy Bush, Gert Doering, Jeff Haas, Jakob Heitz, Nick Hilliard, Igor Malyushkin, Gyan Mishra, Robert Raszuk, John Scudder, and Shyam Sethuram.

Contributors

The following people have made significant contributions to the content of this document.

Avinash Lingala
AT&T
Email: ar977m@att.com

Dhananjaya Rao
Cisco Systems
Email: dhrao@cisco.com

Srihari Sangli
Juniper Networks
Email: ssangli@juniper.net

Authors' Addresses

Jim Uttaro
Individual Contributor
Email: juttaro@ieee.org

Alvaro Retana
Futurewei Technologies, Inc.
Email: alvaro.retana@futurewei.com

Pradosh Mohapatra
Google
Email: pradosh@google.com

Keyur Patel
Arrcus, Inc.
Email: keyur@arrcus.com

Bin Wen
Comcast
Email: bin_wen@comcast.com