

Workgroup: Inter-Domain Routing  
Internet-Draft: draft-uttaro-idr-bgp-oad-01  
Published: 21 April 2023  
Intended Status: Standards Track  
Expires: 23 October 2023  
Authors: J. Uttaro    A. Lingala  
          AT&T            AT&T  
          A. Retana                            P. Mohapatra  
          Futurewei Technologies, Inc.    Google  
          K. Patel            B. Wen        D. Rao  
          Arrcus, Inc.    Comcast    Cisco Systems  
          S. Sangli  
          Juniper Networks

### **One Administrative Domain using BGP**

## **Abstract**

This document defines a new External BGP (EBGP) peering type known as EBGP-OAD. EBGP-OAD peering is used between two EBGP peers 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 23 October 2023.

## **Copyright Notice**

Copyright (c) 2023 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](#)
  - [1.1. Requirements Language](#)
- [2. Discussion](#)
- [3. Operation](#)
  - [3.1. Next Hop Handling](#)
  - [3.2. MULTI\\_EXIT\\_DISC \(MED\) Handling](#)
  - [3.3. Route Reflection](#)
  - [3.4. Tunnel Encapsulation Attribute Handling](#)
  - [3.5. PMSI Tunnel Attribute Handling](#)
  - [3.6. PE Distinguisher Labels Attribute Handling](#)
  - [3.7. BGPsec\\_PATH Attribute Handling](#)
  - [3.8. BGP COMMUNITIES Attribute](#)
  - [3.9. Extended Communities Attribute Handling](#)
  - [3.10. Traffic Engineering Attribute Handling](#)
  - [3.11. IPv6 Address Specific Extended Community Attribute Handling](#)
  - [3.12. Accumulated IGP Metric Attribute Handling](#)
  - [3.13. BGP-LS Attribute Handling](#)
  - [3.14. BGP Role Negotiation](#)
- [4. Deployment and Operational Considerations](#)
- [5. IANA Considerations](#)
- [6. Security Considerations](#)
- [7. References](#)
  - [7.1. Normative References](#)
  - [7.2. Informative References](#)
- [Acknowledgements](#)
- [Authors' Addresses](#)

## 1. Introduction

At each EBGp boundary, BGP path attributes are modified as per standard BGP rules [[RFC4271](#)]. This includes prepending the AS\_PATH attribute with the autonomous-system number of the BGP speaker and stripping any IBGP-only attributes.

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 carry IBGP-only attributes across EBGp peering when the peers belong to OAD. This document defines a new EBGp peering type known as EBGp-OAD. EBGp-OAD peering is used between two EBGp peers that belong to 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 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. An example could be LOCAL\_PREF which is ignored if received on an EBGp session [[RFC4271](#)].

## 3. Operation

[[RFC4271](#)] defines two types of BGP peerings used during a BGP protocol session. As part of the extensions defined in this document, the 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 two external peers in different Autonomous Systems that belong to OAD. In general, the EBGp-OAD speakers follow the EBGp route advertisement, route processing, path attribute announcement and processing rules as defined in [[RFC4271](#)]. In particular, the handling of all attributes with the Transitive bit set should be done in accordance to the announcement and processing rules defined in [[RFC4271](#)].

EBGP-OAD speakers are also allowed to announce and receive any IBGP-only or non-transitive attributes [[RFC4271](#)]. Unless explicitly

specified, all non-transitive path attributes MAY be advertised over an EBGP-OAD session. The reception of any path attribute over an EBGP-OAD session MUST NOT result in an error, unless it is malformed. Received path attributes SHOULD NOT be ignored by the receiver, unless directed to by local policy.

Unless explicitly specified, the current processes for the advertisement of path attributes remains unchanged when advertised through an EBGP-OAD peering. The process for EBGP advertisement MUST take priority over the process for IBGP advertisement. For example, the AS\_PATH attribute is modified as specified in Section 5.1.2 of [\[RFC4271\]](#), bullet b ("BGP speaker advertises the route to an external peer").

An EBGP-OAD speaker MUST support four-octet AS numbers and advertise the "support for four-octet AS number capability" [\[RFC6793\]](#) .

The following sections describe modifications to route advertisements and path attribute announcements that are specific to the EBGP-OAD peering.

### **3.1. Next Hop Handling**

It is reasonable for EBGP-OAD peers to share a common Interior Gateway Protocol (IGP). In such a case, the NEXT\_HOP attribute and the Next Hop in the MP\_REACH\_NLRI attribute [\[RFC4760\]](#) MAY be left unchanged.

### **3.2. MULTI\_EXIT\_DISC (MED) Handling**

The determination of the neighboring AS for the purpose of BGP Route Selection [\[RFC4271\]](#) MAY also consider the ASN of the EBGP-OAD peer. If so, all the peers in the receiving ASN MUST be configured to use the same criteria.

### **3.3. Route Reflection**

BGP Route Reflection [\[RFC4456\]](#) is an alternative to full-mesh IBGP. The ORIGINATOR\_ID and CLUSTER\_LIST attributes MUST NOT be advertised over an EBGP-OAD session. If received, the procedures in [\[RFC7606\]](#) apply.

### **3.4. Tunnel Encapsulation Attribute Handling**

The Tunnel Encapsulation attribute [\[RFC9012\]](#) provides information needed to create tunnels and their corresponding tunnel headers. [\[RFC9012\]](#) restricts the scope of the attribute announcement within a set of ASes that belong to a single administrative entity. Since EBGP-OAD peers belong to a single administrative entity, EBGP-OAD

peers MUST implement the EBGp-related Tunnel Encapsulation attribute procedures defined in [[RFC9012](#)].

### **3.5. PMSI Tunnel Attribute Handling**

The P-Multicast Service Interface Tunnel (PMSI Tunnel) attribute [[RFC6514](#)] provides information needed to create multicast tunnels and their corresponding tunnel headers. EBGp-OAD peers MUST implement the EBGp-related PMSI Tunnel attribute procedures defined in [[RFC6514](#)].

### **3.6. PE Distinguisher Labels Attribute Handling**

PE Distinguisher Labels Attribute attribute [[RFC6514](#)] provides information needed to carry upstream assigned Label values that identifies another PE multicast router. As such, this attribute is only defined to be used inside an AS [[RFC6513](#)]. EBGp-OAD peers MUST implement the procedures defined in [[RFC6513](#)] and [[RFC6514](#)].

### **3.7. BGPsec\_PATH Attribute Handling**

BGPsec\_PATH attribute [[RFC8025](#)] provides security for the path of Autonomous systems through which a BGP UPDATE message passes. EBGp-OAD peers MUST implement the EBGp-related procedures defined in [[RFC8025](#)].

### **3.8. BGP COMMUNITIES Attribute**

BGP COMMUNITIES [[RFC1997](#)] is a transitive attribute used to pass additional information to BGP peers. The advertisement semantics do not change. In particular, routes received carrying the COMMUNITIES attribute containing the well-known NO\_EXPORT value MUST NOT be advertised across an EBGp-OAD session.

### **3.9. Extended Communities Attribute Handling**

The Extended Communities Attribute [[RFC4360](#)] is a transitive attribute that provides a mechanism for labeling information carried in BGP. The Transitive bit is used to indicate whether a particular community is transitive or non-transitive across an Autonomous System (AS) boundary. As described in [[RFC4360](#)], the advertisement of transitive extended communities is subject to local policy for EBGp-OAD peerings. Non-transitive extended communities MAY be advertised to peers over an EBGp-OAD session. For example, the Origin Validation State Extended Community [[RFC8097](#)] can be advertised to peers in the same OAD.

### **3.10. Traffic Engineering Attribute Handling**

The Traffic Engineering attribute is a non-transitive attribute that enables BGP to carry Traffic Engineering information [[RFC5543](#)]. This attribute MAY be advertised to peers over an EBGp-OAD session.

### **3.11. IPv6 Address Specific Extended Community Attribute Handling**

The IPv6 Address Specific Extended Community Attribute is a transitive attribute that carries address information to be used in IPv6-only environments [[RFC5701](#)]. Modeled after Extended Communities [[RFC4360](#)], a particular community in this attribute can be transitive or non-transitive across an Autonomous System (AS) boundary. Non-transitive extended communities MAY be advertised to peers over an EBGp-OAD session.

### **3.12. Accumulated IGP Metric Attribute Handling**

The Accumulated IGP Metric (AIGP) Attribute allows the advertisement of an accumulated internal routing metric across AS boundaries in an "AIGP administrative domain" [[RFC7311](#)]. The AIGP attribute is enabled or disabled on a session by the AIGP\_SESSION configuration item. For EBGp-OAD sessions, the default value of AIGP\_SESSION SHOULD be "enabled".

### **3.13. BGP-LS Attribute Handling**

The BGP Link-State (BGP-LS) attribute is used to carry link, node, and prefix parameters and attributes when distributing link-state and TE information using BGP [[I-D.ietf-idr-rfc7752bis](#)]. This attribute MAY be advertised to peers over an EBGp-OAD session.

### **3.14. BGP Role Negotiation**

Given the intent of OAD, the BGP Role negotiation and OTC Attribute-based procedures specified in [[RFC9234](#)] are NOT RECOMMENDED to be used between peers in an EBGp-OAD session. The OTC attribute, if present, MUST be preserved unchanged through an EBGp-OAD session. The use and negotiation of BGP Roles between EBGp-OAD peers is outside the scope of this document.

## **4. Deployment and Operational Considerations**

For the EBGp-OAD session to operate as expected, both BGP 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, but the BGP session will remain operational.

The default BGP peering type for a session that is across autonomous systems SHOULD be EBGp. BGP implementation SHOULD provide a configuration-time option to enable the EBGp-OAD session type. If the session type is changed once the BGP connection has been established, the BGP speaker MUST readvertise its entire Adj-RIB-Out to its peer. Requesting a route refresh [[RFC7313](#)] is RECOMMENDED.

The requirement that Import and Export Policies exist [[RFC8212](#)] SHOULD be disabled if both peers are configured with the EBGp-OAD session type.

If multiple peerings exist between two autonomous systems that belong to 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.

The consideration of the ASN of the EBGp-OAD peer to determine the neighboring AS for MED comparison [Section 3.2](#) may result in the creation of persistent route oscillations, similar to the Type II Churn described in [[RFC3345](#)]. [[RFC7964](#)] provides solutions and recommendations to address this issue.

## **5. IANA Considerations**

This memo includes no request to IANA.

## **6. Security Considerations**

This extension to BGP does not change the underlying security issues inherent in the existing BGP protocol, such as those described in [[RFC4271](#)] and [[RFC4272](#)].

This document defines a new BGP session type which combines the path attribute propagation rules for EBGp and IBGP peering. Any existing security considerations related to existing path attributes apply to the new EBGp-OAD session type.

By combining the path attribute propagation rules, IBGP information 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 OAD. If misconfigured, the impact is minimal due to the fact that both [[RFC4271](#)] and [[RFC7606](#)] define mechanisms to deal with unexpected path attributes.

## 7. References

### 7.1. Normative References

**[I-D.ietf-idr-rfc7752bis]**

Talaulikar, K., "Distribution of Link-State and Traffic Engineering Information Using BGP", Work in Progress, Internet-Draft, draft-ietf-idr-rfc7752bis-16, 20 February 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-rfc7752bis-16>>.

**[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>>.

**[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>>.

**[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>>.

**[RFC7313]**

Patel, K., Chen, E., and B. Venkatachalapathy, "Enhanced Route Refresh Capability for BGP-4", RFC 7313, DOI 10.17487/RFC7313, July 2014, <<https://www.rfc-editor.org/info/rfc7313>>.

**[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>>.

**[RFC8025]**

Thubert, P., Ed. and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Paging Dispatch", RFC 8025, DOI 10.17487/RFC8025, November 2016, <<https://www.rfc-editor.org/info/rfc8025>>.

**[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>>.

**[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>>.

**[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>>.

## 7.2. 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>>.

## Acknowledgements

TBD

## Authors' Addresses

Jim Uttaro  
AT&T

Email: [ju1738@att.com](mailto:ju1738@att.com)

Avinash Lingala  
AT&T

Email: [ar977m@att.com](mailto:ar977m@att.com)

Alvaro Retana  
Futurewei Technologies, Inc.

Email: [alvaro.retana@futurewei.com](mailto:alvaro.retana@futurewei.com)

Pradosh Mohapatra  
Google

Email: [pradosh@google.com](mailto:pradosh@google.com)

Keyur Patel  
Arrcus, Inc.

Email: [keyur@arrcus.com](mailto:keyur@arrcus.com)

Bin Wen  
Comcast

Email: [bin\\_wen@comcast.com](mailto:bin_wen@comcast.com)

Dhananjaya Rao  
Cisco Systems

Email: [dhrao@cisco.com](mailto:dhrao@cisco.com)

Srihari Sangli  
Juniper Networks

Email: [ssangli@juniper.net](mailto:ssangli@juniper.net)