

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: September 2, 2017

J. Haas  
Juniper Networks, Inc.  
March 1, 2017

Extended Experimental Path Attributes for BGP  
draft-haas-idr-extended-experimental-01

## Abstract

BGP's primary feature extension mechanism, Optional-Transitive Path Attributes, has proven to be a successful mechanism to permit BGP to be extended. In order to ease various issues during the development of new BGP features, this document proposes an extended experimental Path Attribute to carry prototype features.

## Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [RFC2119] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

## 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 <http://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 September 2, 2017.

## Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

## Internet-DraftExtended Experimental Path Attributes for BGP March 2017

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Extended Experimental Path Attribute . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Usage . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Error Handling . . . . .	<a href="#">4</a>
<a href="#">5.</a>	Moving to an Allocated Code Point . . . . .	<a href="#">4</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">5</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">5</a>
<a href="#">8.</a>	References . . . . .	<a href="#">5</a>
<a href="#">8.1.</a>	Normative References . . . . .	<a href="#">5</a>
<a href="#">8.2.</a>	Informative References . . . . .	<a href="#">6</a>
<a href="#">Appendix A.</a>	Comparisons to Other Features . . . . .	<a href="#">6</a>
<a href="#">Appendix B.</a>	Discussion to this Date . . . . .	<a href="#">7</a>
	Author's Address . . . . .	<a href="#">8</a>

[1.](#) Introduction

BGP's [[RFC4271](#)] primary feature extension mechanism, Optional-Transitive Path Attributes, has proven to be a successful mechanism to permit BGP to be extended. It permits implementations to propagate unknown Path Attributes without understanding their contents, so long as they are syntactically valid.

Path Attributes are encoded in BGP UPDATE messages using a single octet code-point. While this code-point space is relatively small, the rate at which new BGP features are introduced has proven to be slow enough that the potential for exhaustion has not been a significant concern more than twenty years into the deployment of BGP-4. This code point space is managed by IANA under the Standards Action policy [[RFC5226](#)], one of the more restrictive policies in IETF's repertoire. Early allocation [[RFC7120](#)] provides some latitude for allocation of these code points compared to the original [RFC 5226](#)

policy, but is reserved for features that are considered appropriately stable.

Development work on the BGP protocol often requires a code point be assigned to a feature in progress. While code point 255 has been

reserved to be Experimental ([[RFC2042](#)]), developers will often face collisions when attempting to do development on more than a single in-progress feature. Once the feature has reached a level of stability, early allocation should be strongly pursued. It may take some time, however, for features to reach that level of stability.

Due to the general difficulty of getting a public code point during the development process, code point "squatting" (use of a code point that has not been officially allocated) is unfortunately common. In many cases, this is done completely internally and has no impact on the Internet. But sometimes accidents happen and pre-release features ship. Prior to the deployment of the Revised BGP Error Handling Procedures [[RFC7606](#)], this could often be disastrous as different features, or different versions of the same feature, collided with each other and were interpreted as syntax errors and caused BGP peering sessions to reset per [RFC 4271](#) error handling procedures. While it is less disastrous for such collisions to happen in terms of stability of the Internet, what's needed is a way for BGP protocol development to proceed with a little more safety.

This document proposes a new BGP Path Attribute, the BGP Extended Experimental Path Attribute. This Attribute is intended to be used solely for BGP Protocol development and is not intended to replace the allocation policies for the BGP Protocol.

## [2.](#) Extended Experimental Path Attribute

The Extended Experimental Path Attribute is an Optional-Transitive Path Attribute with a code of TBD. Its contents are a series of TLVs in the following format:

```
  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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Implementor IANA Private Enterprise Number (4 octets)      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Implementor Feature Code Point Number (4 octets)           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---
```

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Version Number (2 octets)   | Feature Length (2 octets)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Implementor IANA Private Enterprise Number is a Private Enterprise Number (PEN) assigned by IANA. [[IANA-PEN](#)]
- o Feature Code Point Number is a code point space under the control of the holder of the PEN.
- o Version Number is an unsigned number intended to convey the version of the feature covered by the Feature Code Point Number

Haas

Expires September 2, 2017

[Page 3]

Internet-DraftExtended Experimental Path Attributes for BGP March 2017

for the implementor. Implementors are encouraged to sequentially number versions of their feature beginning at 1.

- o Feature Length is the length of the Feature Data.
- o Feature Data will be encoded as a BGP Path Attribute value for the experimental feature.

### [3.](#) Usage

A BGP implementor intending to introduce a new standards oriented Path Attribute will select a code point number for their new Path Attribute and assign an initial Version Number. Whenever the format of the feature needs to change, the Version Number MUST also change. This prevents implementations understanding different versions of a pre-standards feature from improperly parsing the attribute.

BGP Experimental features MUST require explicit configuration to recognize a specific Feature Code Point Number, for a given Version Number, for a given PEN. If such configuration is not present, the TLV MUST be ignored.

BGP Experimental features SHOULD NOT carry more than one Version Number of the same Feature Code Point in a given UPDATE. Implementations are encouraged to strip inconsistent Version Numbered TLVs for a given feature when appropriate. For example, if the BGP speaker is configured to support Version Number 2 of an experimental feature, it may discard all TLVs for the Feature Code Point Number that are not 2.

BGP implementations supporting the Extended Experimental Path

Attribute SHOULD strip this attribute by default on external BGP sessions. Explicit configuration SHOULD be required to permit a given PEN+FCPN+VN tuple into the network.

#### [4.](#) Error Handling

If the Extended Experimental Path Attribute is determined to be syntactically invalid, the Attribute discard behavior from [[RFC7606](#)] MUST be used.

#### [5.](#) Moving to an Allocated Code Point

Once an evolving BGP protocol feature reaches a reasonable level of stability, implementations MUST move to a Path Attribute Code Point allocated using the IETF sanctioned procedures. Implementors that publish their PEN+FCN+VN allocations for a given version of their feature in progress are recommended to publish this binding as part of their allocation request to enable short term backward compatibility with their experimental work.

Haas

Expires September 2, 2017

[Page 4]

---

Internet-DraftExtended Experimental Path Attributes for BGP March 2017

While it is possible for implementations of a new feature to rely on experimental deployment for some time, the procedures noted in [Section 3](#) are intended to discourage this behavior by making inter-domain distribution of the experiment fail by default.

#### [6.](#) Security Considerations

This document does not introduce any new security considerations into the BGP-4 protocol. While the injection of unknown or badly formatted Optional-Transitive Path Attributes has been and remains an issue impacting the stability of the Internet, this proposal doesn't increase exposure to that issue. It is rather expected that this proposal helps remediate the accidental attack surface that incremental BGP protocol work exposes to the Internet at large.

[RFC7606] has mitigated the majority of the issues mentioned in the prior paragraph. See that RFC for further information on the history of the problem.

#### [7.](#) IANA Considerations

This document is primarily about issues related to IANA

Considerations. At some point, IANA will be requested to assign a BGP Path Attribute Code number, referenced as TBD early in the document.

## 8. References

### 8.1. Normative References

- [IANA-PEN] "IANA Private Enterprise Number", <<http://pen.iana.org>>.
- [RFC2042] Manning, B., "Registering New BGP Attribute Types", [RFC 2042](#), DOI 10.17487/RFC2042, January 1997, <<http://www.rfc-editor.org/info/rfc2042>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://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, <<http://www.rfc-editor.org/info/rfc4271>>.

Haas

Expires September 2, 2017

[Page 5]

---

Internet-DraftExtended Experimental Path Attributes for BGP March 2017

- [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, <<http://www.rfc-editor.org/info/rfc7606>>.

### 8.2. Informative References

- [I-D.ietf-idr-bgp-attribute-announcement] Patel, K., Uttaro, J., Decraene, B., Henderickx, W., and J. Haas, "Constrain Attribute announcement within BGP", [draft-ietf-idr-bgp-attribute-announcement-00](#) (work in progress), July 2016.
- [ietf-97-idr-code-point-management-slides] Haas, J., "Code Point Management - IETF 97 Slides",

November 2016,  
<<https://www.ietf.org/proceedings/97/slides/slides-97-idr-code-point-management-02.pdf>>.

- [ietf-97-idr-extended-experimental-path-attribute-slides]  
Haas, J., "Extended Experimental Path Attributes - IETF 97 Slides", November 2016,  
<<https://www.ietf.org/proceedings/97/slides/slides-97-idr-extended-experimental-path-attributes-00.pdf>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), DOI 10.17487/RFC5226, May 2008,  
<<http://www.rfc-editor.org/info/rfc5226>>.
- [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,  
<<http://www.rfc-editor.org/info/rfc6368>>.
- [RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", [BCP 100](#), [RFC 7120](#), DOI 10.17487/RFC7120, January 2014, <<http://www.rfc-editor.org/info/rfc7120>>.

## [Appendix A](#). Comparisons to Other Features

Astute readers will note that this is not the first time BGP Path Attributes have been "tunneled" inside of other Path Attributes. [\[RFC6368\]](#) provided a mechanism by which an entire set of Path Attributes could be tunneled inside of attribute 128 for purposes of transparently passing received BGP Path Attributes in an Internet Layer 3 VPN context from one Customer Edge (CE) router to another.

[RFC6368] suffered from two issues:

1. During its initial development, 4-byte AS numbers were starting to be deployed. This lead to a change in the packet format of the feature to accommodate the 4-byte ASes instead of the previous 2-byte versions.
2. While this feature was intended solely to be used in a VPN context, implementations that did not understand it similarly did

not strip it. This caused the VPN routes to carry attribute 128 in an Internet context after they were delivered to the target CE router.

Due to these two issues, routes containing one version of this feature that "escaped into the wild" eventually to be received by other BGP speakers supporting a different version of the feature. Each version would treat their opposite's encoding as a syntax error. This resulted in BGP peering sessions being reset. This, and other similar issues, was a motivation for [\[RFC6368\]](#).

The second issue noted above is the motivation for [\[I-D.ietf-idr-bgp-attribute-announcement\]](#).

#### [Appendix B](#). Discussion to this Date

This proposal was originally well-received on the IDR mailing list and during its presentation at IETF. Comments included comparison to existing mechanisms in LDP and IS-IS; Hannes Gredler notes that the IS-IS feature is not used.

Another set of comments revolved around the structured format of the PEN+FCN+VN and "why couldn't we simply have a very large first-come, first-served code space". While the author agrees that this would serve a very similar behavior, the author's belief after further consideration is that:

- o Involving IANA, even when the process is very light weight, is part of our existing issue. The Enterprise numbering space permits completely internal management during development of new features.
- o There is no fundamental "burden" of multiple implementors rendezvousing around a common PEN+FCN+VN during interoperability testing. The motivation after such testing should be to request a valid BGP Path Attribute code point using existing IETF procedures.

Another comment was about the possibility of utilizing this mechanism as a long-term private BGP Path Attribute feature. Such behavior may

be a valid use case, however, there remains a need to provide for

automatic filtering of experimental work.

This brings the final comment that both this new Path Attribute and potentially each of the experiments in the Feature Data should be covered by [[I-D.ietf-idr-bgp-attribute-announcement](#)] or something similar. This would include additional procedure to provide for remote filtering of the TLVs defined in this document. Progressing this document, and the use case of long term private Path Attributes as noted in the prior section, should be considered after the attribute-announcement draft receives further feedback.

#### Author's Address

Jeffrey Haas  
Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, CA 94089  
US

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