

InterDomain Routing Group (IDR)  
Internet-Draft  
Updates: [4271](#) (if approved)  
Intended status: Standards Track  
Expires: August 29, 2013

S. Hares  
Huawei Technologies (USA)  
J. Scudder  
Juniper Networks  
February 25, 2013

**Update Attribute Flag Low Bits Clarification  
draft-hares-idr-update-attrib-low-bits-fix-01**

Abstract

This draft provides an update to [RFC 4721](#) to clarify the use of the lower-order four bits of the Attribute flag in the Update message.

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 August 29, 2013.

Copyright Notice

Copyright (c) 2013 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 (<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="#">3</a>
<a href="#">2.</a>	Change to <a href="#">RFC 4271 Section 4.3</a> . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Known BGP Implementation Habits . . . . .	<a href="#">4</a>
<a href="#">4.</a>	IANA Considerations . . . . .	<a href="#">4</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">4</a>
<a href="#">6.</a>	Normative References . . . . .	<a href="#">4</a>
	Authors' Addresses . . . . .	<a href="#">5</a>

## 1. Introduction

[RFC4271] specifies in [section 4.3](#) that that the low order four bits of the the Attribute Flags octet are unused, and MUST be zero when sent. There is a disagreement on what when sent means. This draft specifies the meaning.

The issue has been that one school of thought considers that "when sent" means when originated. Another holds that "when sent" means when originated or propagated. This draft takes the second approach of "when sent" being when originated or propagated.

The real issue is that reserved flags are only useful if there is some hope of someday using them for something. If implementations reset these flags on propagation, then a future revision to the BGP specification which introduces a new flag will not be able to propagate the new attribute flag end to end, since it would be very likely that some well-meaning intermediate router would zero on it. The effort to roll out implementations that transited the new flag would almost certainly be prohibitive.

## 2. Change to [RFC 4271 Section 4.3](#)

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+--+--+--+--+--+--+--+--+--+
|  Attr. Flags  |Attr. Type Code|
+-+--+--+--+--+--+--+--+--+--+

```

Original Text:

The lower-order four bits of the Attribute Flags octet are unused. They MUST be zero when sent and MUST be ignored when received

Corrected Text:

The lower-order four bits of the Attribute Flags octet are unused. They MUST be zero when originated or sent. When received, any MUST be accepted and ignored.

### **3. Known BGP Implementation Habits**

The following are BGP implementation habits regarding the unused flag bits

- o always ignore bits received, and always send zero (originated or propagated);
- o always ignore bits received, always send zero bits (originated), and propagate what was received;
- o if non-zero bits are received, drop the peering session;
- o by special condition (policy) handle set bits or set bits, and propagate;and
- o always sets bits under special conditions, and propagates bits.

The reset of BGP sessions based on non-zero bits has been documented at:

<http://mailman.nanog.org/pipermail/nanog/2012-November/053754.html>

Compliance with this draft, as well as [RFC4271], means that routers should not reset BGP sessions if if non-zero lower bits are received.

### **4. IANA Considerations**

This document includes no request to IANA.

### **5. Security Considerations**

This document has no new security cases.

It clarifies some BGP UPDATE packet flag values and thus may aid in improving BGP security. In particular, it makes it even clearer that routers must not reset a session upon receiving unexpected flag values. Behaving otherwise exposes a router to a denial-of-service attack since a distant party might be able to inject such flag values.

### **6. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

#### Authors' Addresses

Susan Hares  
Huawei Technologies (USA)  
2330 Central Expressway  
Santa Clara, CA 95050  
USA

Email: Susan.Hares@huawei.com

John Scudder  
Juniper Networks  
1194 N. Mathilda Ave  
Sunnyvale, CA 94089  
USA

Email: jgs@juniper.net