

Network Working Group	P. Mohapatra	
Internet-Draft	R. Fernando	
Intended status: Standards Track	Cisco Systems	
Expires: August 28, 2010	February 24, 2010	

[TOC](#)

## **BGP Link Bandwidth Extended Community** **draft-ietf-idr-link-bandwidth-01.txt**

### **Abstract**

This document describes an application of BGP extended communities that allows a router to perform unequal cost load balancing.

### **Status of this Memo**

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 28, 2010.

### **Copyright Notice**

Copyright (c) 2010 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 BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

---

## Table of Contents

- [1.](#) Introduction
    - [1.1.](#) Requirements Language
  - [2.](#) Link Bandwidth Extended Community
  - [3.](#) Deployment Considerations
  - [4.](#) Acknowledgments
  - [5.](#) IANA Considerations
  - [6.](#) Security Considerations
  - [7.](#) Normative References
  - [8.](#) Authors' Addresses
- 

## 1. Introduction

[TOC](#)

When a BGP speaker receives multiple paths from its internal peers, it could select more than one path to send traffic to. In doing so, it might be useful to provide the speaker with information that would help it distribute the traffic unequally based on the cost of the external (DMZ) link. This document suggests that the external link bandwidth be carried in the network using a new extended community [\[RFC4360\]](#) (Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute," February 2006.) - the link bandwidth extended community.

---

### 1.1. Requirements Language

[TOC](#)

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

---

## 2. Link Bandwidth Extended Community

[TOC](#)

When a BGP speaker receives a route from a directly connected external neighbor (the external neighbor that is one IP hop away) and advertises this route (via IBGP) to internal neighbors, as part of this advertisement the router may carry the bandwidth of the link that connects the router with the external neighbor. The bandwidth of such a link is carried in the Link Bandwidth Community. The community is optional non-transitive. A border router **MUST** strip the link bandwidth community from a route when it advertises the route to an external neighbor. The value of the high-order octet of the extended Type Field is 0x40. The value of the low-order octet of the extended type field for this community is 0x04. The value of the Global Administrator subfield in the Value Field **SHOULD** represent the Autonomous System of the router that attaches the Link Bandwidth Community. If four octet AS numbering scheme is used [\[RFC4893\] \(Vohra, Q. and E. Chen, "BGP Support for Four-octet AS Number Space," May 2007.\)](#), AS\_TRANS should be used in the Global Administrator subfield. The bandwidth of the link is expressed as 4 octets in IEEE floating point format, units being bytes per second. It is carried in the Local Administrator subfield of the Value Field.

---

## 3. Deployment Considerations

[TOC](#)

The usage of this community is restricted to the cases where BGP multipath can be safely deployed. In other words, the IGP distance between the load balancing router and the exit points should be the same. Alternatively, the path between the load sharing router and the exit points could be label switched. If there are multiple paths to reach a destination and if only some of them have link bandwidth community, the receiver should not perform unequal cost load balancing based on link bandwidths.

---

## 4. Acknowledgments

[TOC](#)

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.

---

## 5. IANA Considerations

[TOC](#)

This document defines a specific application of the two-octet AS specific extended community. IANA is requested to assign a sub- type value of 0x04 for the link bandwidth extended community.

Name	Value
----	-----
non-transitive Link Bandwidth Ext. Community	0x4004

---

## 6. Security Considerations

[TOC](#)

There are no additional security risks introduced by this design.

---

## 7. Normative References

[TOC](#)

[RFC2119]	<a href="#">Bradner, S.</a> , " <a href="#">Key words for use in RFCs to Indicate Requirement Levels</a> ," BCP 14, RFC 2119, March 1997 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).
[RFC4360]	<a href="#">Sangli, S.</a> , <a href="#">Tappan, D.</a> , and <a href="#">Y. Rekhter</a> , " <a href="#">BGP Extended Communities Attribute</a> ," RFC 4360, February 2006 ( <a href="#">TXT</a> ).
[RFC4893]	<a href="#">Vohra, Q.</a> and <a href="#">E. Chen</a> , " <a href="#">BGP Support for Four-octet AS Number Space</a> ," RFC 4893, May 2007 ( <a href="#">TXT</a> ).

---

## Authors' Addresses

[TOC](#)

	Pradosh Mohapatra
	Cisco Systems
	170 W. Tasman Drive
	San Jose, CA 95134
	USA
Phone:	
Email:	<a href="mailto:pmohapat@cisco.com">pmohapat@cisco.com</a>
	Rex Fernando
	Cisco Systems
	170 W. Tasman Drive
	San Jose, CA 95134
	USA

Phone:	
Email:	<a href="mailto:rex@cisco.com">rex@cisco.com</a>