

IDR Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 19, 2011

R. Raszuk  
Cisco Systems  
J. Haas  
Juniper Networks  
October 16, 2010

Registered Wide BGP Community Values  
draft-raszuk-registered-wide-bgp-community-values-00

## Abstract

Communicating various routing policies via route tagging plays an important role in external BGP peering relations. The most common tool used today to attach various information about routes is realized with the use of BGP communities. Such information is important for the peering AS to perform some mutually agreed actions without the need to maintain a separate offline database for each pair of prefix and an associated with it requested set of action entries.

This document proposes to establish a new IANA maintained registry of most commonly used Wide BGP Communities by network operators. Such public registry will allow for easy reference and clear interpretation of the actions associated with received community values.

## 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 April 19, 2011.

## Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft registered-wide-bgp-community-values

October 2010

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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>	<a href="#">Introduction . . . . .</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Globally significant pre-defined values . . . . .</a>	<a href="#">3</a>
<a href="#">2.1.</a>	<a href="#">Well Known Standard BGP Communities . . . . .</a>	<a href="#">3</a>
<a href="#">2.2.</a>	<a href="#">Registered pre-defined Wide BGP Communities . . . . .</a>	<a href="#">3</a>
<a href="#">2.2.1.</a>	<a href="#">General Registered Wide BGP Community Values . . . . .</a>	<a href="#">4</a>
<a href="#">2.2.2.</a>	<a href="#">Advertisement control Registered Wide BGP Communities . . . . .</a>	<a href="#">7</a>
<a href="#">2.2.3.</a>	<a href="#">AS source marking Registered Wide BGP Communities . . . . .</a>	<a href="#">8</a>
<a href="#">2.2.4.</a>	<a href="#">Return path influencing Registered Wide BGP Communities . . . . .</a>	<a href="#">10</a>
<a href="#">2.2.5.</a>	<a href="#">AS_PATH modifying Registered Wide BGP Communities . . . . .</a>	<a href="#">10</a>
<a href="#">2.2.6.</a>	<a href="#">Local Preference Registered Community . . . . .</a>	<a href="#">11</a>
<a href="#">2.2.7.</a>	<a href="#">AS_PATH TTL Registered Community . . . . .</a>	<a href="#">12</a>
<a href="#">3.</a>	<a href="#">Example . . . . .</a>	<a href="#">12</a>
<a href="#">4.</a>	<a href="#">Security considerations . . . . .</a>	<a href="#">13</a>
<a href="#">5.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">13</a>
<a href="#">6.</a>	<a href="#">Contributors . . . . .</a>	<a href="#">14</a>
<a href="#">7.</a>	<a href="#">Acknowledgments . . . . .</a>	<a href="#">15</a>
<a href="#">8.</a>	<a href="#">References . . . . .</a>	<a href="#">16</a>
<a href="#">8.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">16</a>
<a href="#">8.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">16</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">16</a>

---

Internet-Draft      registered-wide-bgp-community-values      October 2010

## [1.](#) Introduction

[RFC 1997](#) [[RFC1997](#)] defines a BGP Community Attribute to be used as a tool to contain in BGP update message various additional information about routes which may help to automate peering administration. As defined in [RFC 1997](#) [[RFC1997](#)] BGP Communities attribute consists of one or more sets of four octet values, where each one of them specifies a different community. Except two reserved ranges the encoding of community values mandates that first two octets are to contain the Autonomous System number followed by next two octets containing locally defined value.

This document lists the most commonly used today BGP communities as well as provides a new registry for future definitions.

## [2.](#) Globally significant pre-defined values

### [2.1.](#) Well Known Standard BGP Communities

According to [RFC 1997](#) as well as to IANA's Well-Known BGP Communities registry today the following BGP communities are defined to have global significance:

0xFFFF0000	planned-shut	<a href="#">[draft-francois-bgp-gshut]</a>
0xFFFFFFFF01	NO_EXPORT	<a href="#">[RFC1997]</a>
0xFFFFFFFF02	NO_ADVERTISE	<a href="#">[RFC1997]</a>
0xFFFFFFFF03	NO_EXPORT_SUBCONFED	<a href="#">[RFC1997]</a>
0xFFFFFFFF04	NOPEER	<a href="#">[RFC3765]</a>

This document recommends for simplicity as well as for avoidance of backward compatibility issues the continued use of BGP Standard Community Attribute type 8 as defined in [RFC 1997](#) to distribute non Autonomous System specific Well-Known BGP Communities.

For the same reason the described registry does not intend to obsolete BGP Extended Community Attribute and any already defined and deployed extended communities. The new registry is to be used primarily for new community definitions in particular those which require to carry various new parameters or which should be propagated with a controlled scope and radius.

## [2.2.](#) Registered pre-defined Wide BGP Communities

It has been requested numerous times to have a globally unified way to express some particular Autonomous System based routing policies.

When defining a new way to encode bgp communities we have an opportunity to define set of new registered routing policies and route markings which could be passed within and between Autonomous Systems resulting in their common interpretation.

This document will request IANA to define and maintain a new registry for pre-defined Wide BGP Community values. The allocation policy is on a first come first served basis.

It is recommended that an implementation supports by an explicit enabling defined below Registered Wide BGP Communities. Depending on the BGP implementation support it is recommended that an implementation would support Registered Wide BGP Communities without breaking static or dynamic peer/update groups. However it needs to be pointed out that support of all Registered Wide BGP Communities is not mandatory. It will be perfectly valid for any BGP implementation to support only subset of Wide BGP Communities.

It is strongly advised that each Autonomous System does an inbound verification of received Wide BGP Communities from all of its EBGP peers before accepting them and propagating within their own domain.

The document does not mandate nor enforces that given registered type value of Wide BGP Community would be of transitive or non-transitive type. It is for the operator to determine the propagation AS radius required for such community when appending it to routing information. However the document will provide a transitivity radius recommendation to defined communities.

The following Wide BGP Communities have global significance and their execution should be uniformly implemented by any BGP speaker supporting given set of Wide BGP Communities.

The defined below value of the community should be interpreted as registered value only if "R" - registered bit is set in the community Type 1 container as described in [[draft-raszuk-wide-bgp-communities](#)] Otherwise the value is local and it's actions is locally defined by the operator.

### [2.2.1.](#) General Registered Wide BGP Community Values

The below set of communities will be defined to be carried in Wide BGP Community Type 1, with the container type values as per [Section 5](#).

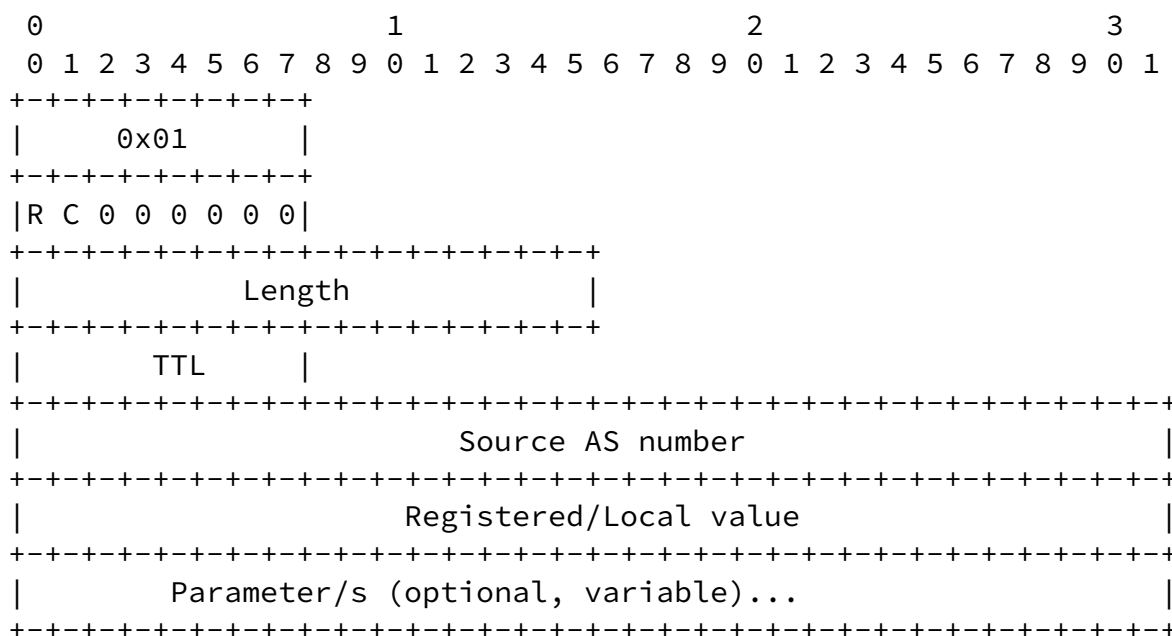


Figure 1: Wide BGP Community Type 1

Description format:

## NAME

TTL - defines domain or sub-domain propagation radius  
LENGTH - Length of the TLV in octets  
PARAMETERS - optional parameters for execution  
DESCRIPTION - defines the use and post execution action

## BLACKHOLE

TTL - Operator's defined  
LENGTH - 13 octets  
PARAMETERS - none

DESCRIPTION - All transit traffic to destinations for which advertised routes carry such community value should be dropped. It is recommended that specified Autonomous System number should be eligible and verified by BGP Origin Validation functionality to advertise given BGP destinations.

## SOURCE BLACKHOLE

TTL - Operator's defined  
LENGTH - 13 octets  
PARAMETERS - none

DESCRIPTION - All transit traffic which source addresses have been tagged by such Wide BGP Community should be dropped.

## SOURCE DO RPF

TTL - Operator's defined  
LENGTH - 13 octets  
PARAMETERS - none

DESCRIPTION - All transit traffic which source addresses have been tagged by such Wide BGP Community should be subject to Reverse Path Forwarding check when crossing Autonomous System boundaries. Source Autonomous System number specified in the body of this community should directly indicate the peering interfaces on which such RPF check should be performed.

#### HIGH PRIORITY PREFIX

TTL - 0x00 (local to domain)

LENGTH - 13 octets

PARAMETERS - none

DESCRIPTION - BGP prefixes carrying such Wide BGP Community should be advertised to restarting peers before other prefixes received by given BGP speaker.

#### ATTACK TARGET

TTL - Operator's defined

LENGTH - 13 octets

PARAMETERS - none

DESCRIPTION - The ATTACK\_TARGET Registered Wide BGP Community indicates that BGP prefixes carrying such community are receiving unusual amount of unwanted traffic most likely due to some form of network attack. Network devices capable of analyzing and mitigating such attacks can use such community as a hint on what destinations to focus the most.

#### [2.2.2.](#) Advertisement control Registered Wide BGP Communities

#### NO ADVERTISE TO AS

TTL - Operator's defined

LENGTH - 17 octets

PARAMETERS - Target AS number - 4 octets

DESCRIPTION - All routes received which carry such Wide BGP Community containing this value MUST NOT be advertised to BGP peer which Autonomous System number has been listed in the body of this community.

Semantically specifying the reserved Autonomous System value of 0xFFFFFFFF (ANY AS) would be an equivalent of using NO\_ADVERTISE Well-Known Standard BGP Community Attribute.

#### ADVERTISE TO AS

TTL - Operator's defined

LENGTH - 17 octets

PARAMETERS - Target AS number - 4 octets

DESCRIPTION - All routes received carrying such Wide BGP Community containing this value MUST ONLY be advertised to BGP peers which Autonomous System number is specified in the PARAMETER field of this community.

Semantically specifying the reserved Autonomous System value of 0xFFFFFFFF (ANY AS) would be an equivalent of advertisement to all neighbors. Post execution this community MUST be removed.

#### ADVERTISE AND SET NO EXPORT

TTL - Operator's defined

LENGTH - 17 octets

PARAMETERS - Target AS number - 4 octets

DESCRIPTION - All routes received carrying such Wide BGP Community containing this value MUST be advertised to BGP peer which Autonomous System number is specified in the PARAMETER section of this community with NO\_EXPORT Standard BGP Community attached.

Semantically specifying the reserved Autonomous System value of 0xFFFFFFFF (ANY AS) would be an equivalent of advertisement to all

neighbors with NO\_EXPORT community being set. Post execution this



community MUST be removed.

### [2.2.3.](#) AS source marking Registered Wide BGP Communities

#### FROM PEER

TTL - 0x00 - Local AS tagging only  
LENGTH - 13 octets  
PARAMETERS - none

DESCRIPTION - Autonomous System may attach this community to routes received from their EBGP peers to later, when advertising them outside the domain, apply or relax local policies only on such group of destinations.

#### FROM CUSTOMER

TTL - 0x00 - Local AS tagging only  
LENGTH - 13 octets  
PARAMETERS - none

DESCRIPTION - Autonomous System may attach this community to routes received from their customers to later, when advertising them outside the domain, apply or relax local policies only on such group of destinations.

#### INTERNAL

TTL - 0x00 - Local AS tagging only  
LENGTH - 13 octets  
PARAMETERS - none

DESCRIPTION - Autonomous System may attach this community to routes originated in their own domain to later, when advertising them outside the domain, apply or relax local policies only on such group of destinations.

## FROM UPSTREAM

TTL - 0x00 - Local AS tagging only  
LENGTH - 13 octets  
PARAMETERS - none

DESCRIPTION - Autonomous System may attach this community to routes received from their EBGp upstream peers to later, when advertising them outside the domain, apply or relax local policies only on such group of destinations.

## FROM IX

TTL - 0x00 Local AS tagging only  
LENGTH - 13 octets  
PARAMETERS - none

DESCRIPTION - Autonomous System may attach this community to routes received from their EBGp peering sessions with the Internet Exchange peers or with Route Server to later, when advertising them outside the domain, apply or relax local policies only on such group of destinations.

## LEARNED FROM AS

TTL - 0x00 Local AS tagging only  
LENGTH - 17 octets  
PARAMETERS - 4 octets - Peer's AS number

DESCRIPTION - Autonomous System may attach this community to routes received from their EBGp peer by explicitly tagging them with their peer's Autonomous System number as a value of the PARAMETER field. If the AS number is a two octet number first two octet will be filled with zero. It is possible to use this to also carry private AS number of customers.

Internet-Draft

registered-wide-bgp-community-values

October 2010

#### [2.2.4.](#) Return path influencing Registered Wide BGP Communities

##### PATH HINT

TTL - Operator's defined

LENGTH - 17 octets

PARAMETERS - 4 octets AS number

DESCRIPTION - Autonomous System receiving such Wide BGP Community value should prefer for BGP prefixes received with such community (for example by increasing value of local preference on ingress), a BGP path which traverses Autonomous System number which has been specified in the PARAMETER field of this community. Post execution this community SHOULD be kept.

##### NEGATIVE PATH HINT

TTL - Operator's defined

LENGTH - 17 octets

PARAMETERS - 4 octets AS number

DESCRIPTION - Autonomous System receiving such Wide BGP Community value should prefer for BGP prefixes received with such community (for example by increasing value of local preference on ingress), a BGP path which DOES NOT traverses Autonomous System number which has been specified in the PARAMETERS field of this community. Post execution this community SHOULD be kept.

#### [2.2.5.](#) AS\_PATH modifying Registered Wide BGP Communities

##### PREPEND N TIMES BY AS

TTL - Operator's defined

LENGTH - 18 octets

PARAMETERS - 4 octets AS number + 1 octet prepend's number

DESCRIPTION - The Autonomous System specified in the PARAMETERS field of such community should prepend N times its own Autonomous System number when advertising routes tagged with this community to peers. Number of requested AS prepends is provided in the last octet of PARAMETERS field value. Post execution this community MUST be removed.

Internet-Draft      registered-wide-bgp-community-values      October 2010

#### PREPEND N TIMES TO AS

TTL - Operator's defined

LENGTH - 18 octets

PARAMETERS - 4 octets AS number + 1 octet prepend's number

DESCRIPTION - The Autonomous System advertising routes externally should prepend N times its own Autonomous System number when advertising routes tagged with this community to peer which AS number is defined by the first 4 octets of PARAMETERS field. Number of requested AS prepends is provided in the last octet of PARAMETERS field value. Post execution this community MUST be removed.

#### REPLACE BY

TTL - Operator's defined

LENGTH - 17 octets

PARAMETERS - 4 octets AS number

DESCRIPTION - All routes marked with such community advertised by an Autonomous System to all of its external peers should have any occurrence of an Autonomous System number specified in the PARAMETERS field replaced with advertising domain's local Autonomous System number. Post execution this community MUST be removed.

### [2.2.6.](#) Local Preference Registered Community

#### LOCAL PREFERENCE

TTL - Operator's defined

LENGTH - 14 octets

PARAMETERS - 1 octet 1st bit indicates:

0-increment, 1-decrement;

7 bits - value of local preference value 1..127

DESCRIPTION - Autonomous System may suggest to its EBGp neighbor the following adjustments to the value of local preference as specified by given domain's local policy. The values of requested increment or decrement local preference value is carried in the PARAMETERS field. Post execution this community MUST be removed.

#### [2.2.7.](#) AS\_PATH TTL Registered Community

##### AS\_PATH TTL MAX RADIUS

TTL - Operator's defined

LENGTH - 14 octets

PARAMETERS - 1 octet - max AS\_PATH radius

DESCRIPTION - Autonomous System may suggest to drop advertised prefix by any transit network if its AS\_PATH attribute length would be equal or greater to encoded value both inbound or outbound of EBGp session. The value of max AS\_PATH length allowed is specified in the PARAMETERS field of the community. Post comparison this community MUST be kept.

### [3.](#) Example

Customer of the source AS number XYZ requests to execute AS\_PATH prepend 4 times when advertising the prefixes to AS number 2424. We will use the following community assigned on ingress or at the prefix origination.

```
TTL - 0x00
LENGTH - 18 octets
VALUE - 17 / 0x12
PARAMETERS - 4 octets AS number + 1 octet prepend's number
```

[illegible]

+--+--+--+--+--+--+

#### 4. Security considerations

All the security considerations for BGP Communities as well as for BGP Extended Communities RFCs apply here.

#### 5. IANA Considerations

This document requests IANA to define and maintain a new registry named: "Registered Wide BGP Communities Values". The reserved pool of 0x00000000-0xFFFFFFFF has been defined for its allocations. The allocation policy is on a first come first served basis.

This document makes the following assignments for the Registered Wide BGP Community values:

+-----+-----+	
Name	Type Value
+-----+-----+	
BLACKHOLE	1
SOURCE BLACKHOLE	2
SOURCE DO RPF	3
HIGH PRIORITY PREFIX	4
ATTACK TARGET	5
NO ADVERTISE TO AS	6
ADVERTISE TO AS	7
ADVERTISE AND SET NO EXPORT	8

FROM PEER	9	
FROM CUSTOMER	10	
INTERNAL	11	
FROM UPSTREAM	12	
FROM IX	13	
LEARNED FROM AS	14	
PATH HINT	15	
PATH NEGATIVE HINT	16	
PREPEND N TIMES BY AS	17	
PREPEND N TIMES TO AS	18	
REPLACE BY	19	
LOCAL PREFERENCE	20	
AS_PATH TTL MAX RADIUS	21	
FREE POOL	22..	
+-----+-----+		

## 6. Contributors

The following people contributed significantly to the content of the document:

Bruno Decraene  
France Telecom  
38-40 rue du General Leclerc  
92794 Issi Moulineaux cedex 9  
France  
Email: [bruno.decraene@orange-ftgroup.com](mailto:bruno.decraene@orange-ftgroup.com)

Shintaro Kojima  
OTEMACHI 1st. SQUARE EAST TOWER, 3F  
1-5-1, Otemachi,  
Chiyoda-ku, Tokyo 100-0004



Japan  
Email: koji@mfeed.ad.jp

Juan Alcaide  
Cisco Systems  
Research Triangle Park, NC  
United States  
Email: jalcaide@cisco.com

Burjiz Pithawala  
Cisco Systems  
170 West Tasman Dr  
San Jose, CA  
United States  
Email: bpithaw@cisco.com

Saku Ytti  
TDC Oy  
Mechelininkatu 1a  
00094 TDC  
Finland  
Email: ytti@tdc.net

Paul Jakma  
School of Computing Science, Uni. of Glasgow  
Sir Alwyn Williams Building  
University of Glasgow  
Glasgow  
G1 5AE  
UK  
Email: paulj@dcs.gla.ac.uk

## 7. Acknowledgments

Authors would like to thank Enke Chen, Pedro Marques and Alton Lo for their valuable input.

## 8. References

### 8.1. 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.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), February 2006.

### 8.2. Informative References

- [RFC1997] Chandrasekeran, R., Traina, P., and T. Li, "BGP Communities Attribute", [RFC 1997](#), August 1996.
- [RFC1998] Chen, E. and T. Bates, "An Application of the BGP Community Attribute in Multi-home Routing", [RFC 1998](#), August 1996.
- [RFC4384] Meyer, D., "BGP Communities for Data Collection", [BCP 114](#), [RFC 4384](#), February 2006.
- [RFC4893] Vohra, Q. and E. Chen, "BGP Support for Four-octet AS Number Space", [RFC 4893](#), May 2007.
- [RFC5668] Rekhter, Y., Sangli, S., and D. Tappan, "4-Octet AS Specific BGP Extended Community", [RFC 5668](#), October 2009.

### Authors' Addresses

Robert Raszuk  
Cisco Systems  
170 West Tasman Drive  
San Jose, CA 95134  
US

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

Internet-Draft      registered-wide-bgp-community-values

October 2010

Jeffrey Haas  
Juniper Networks  
1194 N.Mathilda Ave  
Sunnyvale, CA 94089  
US

Email: [jhaas@pfrc.org](mailto:jhaas@pfrc.org)

Raszuk & Haas

Expires April 19, 2011

[Page 17]