

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2015

R. Raszuk, Ed.
Mirantis Inc.
J. Haas
Juniper Networks
March 7, 2015

Registered Wide BGP Community Values
draft-raszuk-registered-wide-bgp-communities-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 September 8, 2015.

Copyright Notice

Copyright (c) 2015 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

1. Introduction	2
2. Globally significant pre-defined values	3
2.1. Well Known Standard BGP Communities	3
2.2. Registered pre-defined Wide BGP Communities	3
2.2.1. General Registered Wide BGP Community Values	4
2.2.2. Advertisement control Registered Wide BGP Communities	7
2.2.3. AS source marking Registered Wide BGP Communities	8
2.2.4. Return path influencing Registered Wide BGP Communities	10
2.2.5. AS_PATH modifying Registered Wide BGP Communities	10
2.2.6. Local Preference Registered Community	11
2.2.7. AS_PATH TTL Registered Community	12
2.2.8. GEO-LOCATION Registered Community	12
3. Example	13
4. Security considerations	14
5. IANA Considerations	14
6. Contributors	15
7. Acknowledgments	16
8. References	16
8.1. Normative References	16
8.2. Informative References	17
Authors' Addresses	17

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	[draft-francois-bgp-gshut]
0FFFFFFF01	NO_EXPORT	[RFC1997]
0FFFFFFF02	NO_ADVERTISE	[RFC1997]
0FFFFFFF03	NO_EXPORT_SUBCONFED	[RFC1997]
0FFFFFFF04	NOPEER	[RFC3765]

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 (Community Registered Value) as per Section 5.

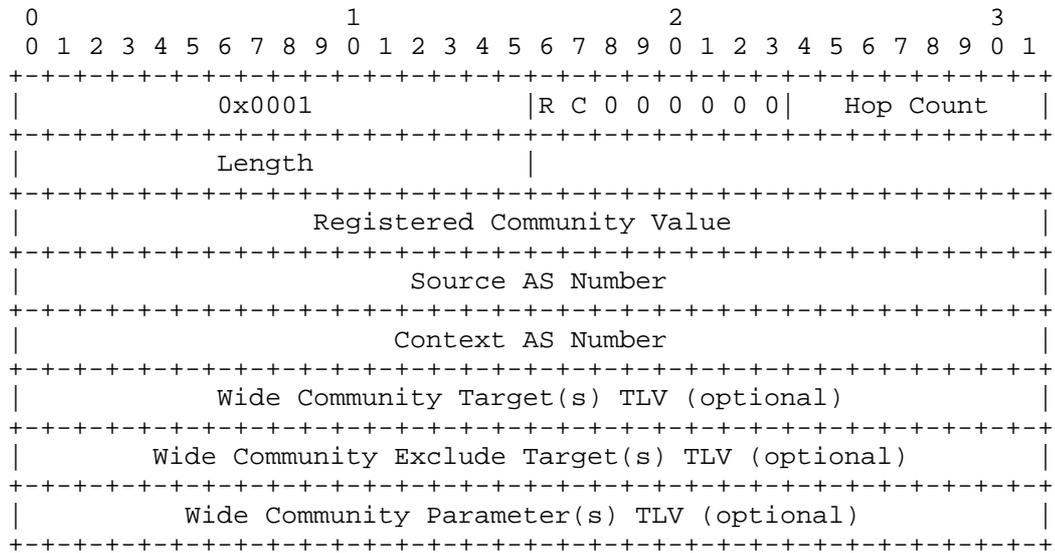


Figure 1: Wide BGP Community Type 1

Description format:

- TYPE:
 - 0x0001 (constant for this registry)
- FLAGS "F":
 - R - Registered bit (Set to 1 for registered values)
 - C - Confederation bit (Set when applicable)
- HOP COUNT "H":
 - Defines domain or sub-domain propagation radius
- LENGTH "L":
 - Length of the Container Type 1 in octets
- REGISTERED COMMUNITY VALUE "R":
 - Value of the community in registry
- SOURCE AS "S":
 - Originator AS of Wide BGP Community
- CONTEXT AS "C":
 - For registered communities carries predefined meaning or otherwise should be set to 0x00000000
- TARGET TLV "T":
 - Set of atoms containing targets for execution
- EXCLUDE TARGET TLV "E":
 - Set of atoms containing excluded targets for execution
- PARAMETER TLV "P":
 - Set of atoms containing optional parameters for execution

BLACKHOLE

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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 FILTER

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = none
L = 18 octets	E = none
R = IANA assigned	P = none

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

SOURCE DO RPF

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = 0x00	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = Type_1 (Peer_AS)
L = 25 octets	E = none
R = IANA assigned	P = none

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 TARGET TLV field 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = Type_1 (Peer_AS)
L = 25 octets	E = none
R = IANA assigned	P = none

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 TARGET TLV 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = Type_1 (Peer_AS)
L = 25 octets	E = none
R = IANA assigned	P = none

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 TARGET TLV field of this community with NO_EXPORT Standard BGP Community attached.

Semantically specifying in TARGET TLV 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = 0x00	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = 0x00	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = 0x00	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = 0x00	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = 0x00	T = none
L = 18 octets	E = none
R = IANA assigned	P = 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = 0x00	T = Type_1 (Peer_AS)
L = 25 octets	E = none
R = IANA assigned	P = none

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 TARGET TLV 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.

2.2.4. Return path influencing Registered Wide BGP Communities

PATH HINT

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = Type_1 (AS#)
L = 25 octets	E = none
R = IANA assigned	P = none

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 TARGET TLV field of this community. Post execution this community SHOULD be kept.

NEGATIVE PATH HINT

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = Type_1 (AS#)
L = 25 octets	E = none
R = IANA assigned	P = none

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 TARGET TLV 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = Type_1 (AS#)
L = 29 octets	E = none
R = IANA assigned	P = Type_4 (0xAA)

DESCRIPTION - The Autonomous System specified in the TARGET TLV field of such community should prepend N times (encoded as 0xAA) its own Autonomous System number when advertising routes tagged with this community to peers. Number of requested AS prepends is provided in the PARAMETERS TLV field value. Post execution this community MUST be removed.

PREPEND N TIMES TO AS

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = Type_1 (AS#)
L = 29 octets	E = none
R = IANA assigned	P = Type_4 (0xAA)

DESCRIPTION - The Autonomous System advertising routes externally should prepend N times (encoded as 0xAA) its own Autonomous System number when advertising routes tagged with this community to peer which AS number is defined by TARGET TLV field. Number of requested AS prepends is provided in the PARAMETERS TLV field. Post execution this community MUST be removed.

REPLACE BY

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = Type_1 (AS#)
L = 25 octets	E = none
R = IANA assigned	P = none

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 TARGET TLV 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

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = none
L = 22 octets	E = none
R = IANA assigned	P = Type_4 (ABBBBBBB)

SEMANTICS of PARAMETERS TLV

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 of local preference value is carried in the PARAMETERS TLV field. Post execution this community MUST be removed.

2.2.7. AS_PATH TTL Registered Community

AS_PATH TTL MAX RADIUS

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = none
L = 22 octets	E = none
R = IANA assigned	P = Type_4 (0xAA) 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 TLV field of the community. Post comparison this community MUST be kept.

2.2.8. GEO-LOCATION Registered Community

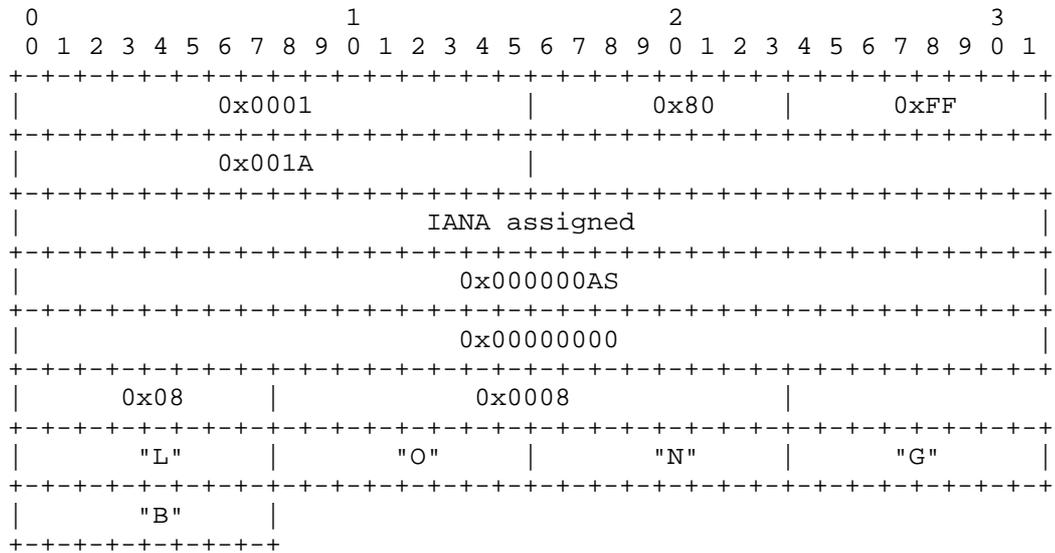
GEOGRAPHIC LOCATION WHERE BGP ROUTE IS INTRODUCED TO AS

Type: 0x0001	S = src AS #
F = 0x80	C = 0x00000000
H = Operator's defined	T = none
L = 26 octets	E = none
R = IANA assigned	P = Type_8 (5 UTF-8 characters)

DESCRIPTION - Autonomous Systems may attach this community to routes received from EBGp neighbors or introduced to BGP by other routing protocols to identify the geographic location where the route was introduced to the AS. The "right-most" two octets of PARAMETERS TLV correspond to an ISO3166-1 alpha-2 country identifier, while the "left-most" three octets may express a more specific geographic location, such as a city or IXP encoded in 3 octets.

Example:

Wide BGP Community describing route learnt by the AS at London, GB
 HOP COUNT - operator defined
 LENGTH - 26
 PARAMETERS - 3 octets locality string + 2 octets country id.



3. Example

Customer of the source AS number 100 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.

```

PREPEND N TIMES TO AS
  Type: 0x0001           S = 0x00000064 (dec 100)
  F = 0x80              C = 0x00000000
  H = 0x00              T = 0x00000978 (dec 2424)
  L = 0x001D (dec 29 octets) E = none
  R = IANA assigned     P = 0x04 (dec 4)

```

0										1										2										3									
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	2	3	4	5	6	7	8	9
0x0001										0x80										0x00																			
0x001D																																							
IANA assigned																																							
0x00000064																																							
0x00000000																																							
0x01										0x0007																													
0x00000978																																							
0x04										0x0004																													
0x04																																							

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 FILTER	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
GEO-LOCATION	22
FREE POOL	23..

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

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, Alton Lo and Jeff Wheeler 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 (editor)
Mirantis Inc.
615 National Ave. #100
Mt View, CA 94043
USA

Email: robert@raszuk.net

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

Email: jhaas@pfrc.org