

Network Working Group
Request for Comments: 4360
Category: Standards Track

S. Sangli
D. Tappan
Cisco Systems
Y. Rekhter
Juniper Networks
February 2006

BGP Extended Communities Attribute

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes the "extended community" BGP-4 attribute. This attribute provides a mechanism for labeling information carried in BGP-4. These labels can be used to control the distribution of this information, or for other applications.

1. Introduction

The Extended Community Attribute provides a mechanism for labeling information carried in BGP-4 [BGP-4]. It provides two important enhancements over the existing BGP Community Attribute [RFC1997]:

- An extended range, ensuring that communities can be assigned for a plethora of uses, without fear of overlap.
- The addition of a Type field provides structure for the community space.

The addition of structure allows the usage of policy based on the application for which the community value will be used. For example, one can filter out all communities of a particular type, or allow only certain values for a particular type of community. It also allows one to specify whether a particular community is transitive or non-transitive across an Autonomous System (AS) boundary. Without structure, this can only be accomplished by explicitly enumerating

all community values that will be denied or allowed and passed to BGP speakers in neighboring ASes based on the transitive property.

1.1. Specification of Requirements


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 [RFC2119].

2. BGP Extended Communities Attribute

The Extended Communities Attribute is a transitive optional BGP attribute, with the Type Code 16. The attribute consists of a set of "extended communities". All routes with the Extended Communities attribute belong to the communities listed in the attribute.

Each Extended Community is encoded as an 8-octet quantity, as follows:

- Type Field : 1 or 2 octets
- Value Field : Remaining octets



(*) Present for Extended types only, used for the Value field otherwise.

Type Field:

Two classes of Type Field are introduced: Regular type and Extended type.

The size of Type Field for Regular types is 1 octet, and the size of the Type Field for Extended types is 2 octets.

The value of the high-order octet of the Type Field determines if an extended community is a Regular type or an Extended type. The class of a type (Regular or Extended) is not encoded in the structure of the type itself. The class of a type is specified in the document that defines the type and the IANA registry.

The high-order octet of the Type Field is as shown below:

```

  0 1 2 3 4 5 6 7
  +-----+
  |I|T|         |
  +-----+

```

I - IANA authority bit

Value 0: IANA-assignable type using the "First Come First Serve" policy

Value 1: Part of this Type Field space is for IANA assignable types using either the Standard Action or the Early IANA Allocation policy. The rest of this Type Field space is for Experimental use.

T - Transitive bit

Value 0: The community is transitive across ASes

Value 1: The community is non-transitive across ASes

Remaining 6 bits: Indicates the structure of the community

Value Field:

The encoding of the Value Field is dependent on the "type" of the community as specified by the Type Field.

Two extended communities are declared equal only when all 8 octets of the community are equal.


The two members in the tuple <Type, Value> should be enumerated to specify any community value. The remaining octets of the community interpreted based on the value of the Type field.

3. Defined BGP Extended Community Types

This section introduces a few extended types and defines the format of the Value Field for those types. The types introduced here provide "templates", where each template is identified by the high-order octet of the extended community Type field, and the lower-order octet (sub-type) is used to indicate a particular type of extended community.

3.1. Two-Octet AS Specific Extended Community

This is an extended type with Type Field composed of 2 octets and Value Field composed of 6 octets.



The value of the high-order octet of this extended type is either 0x00 or 0x40. The low-order octet of this extended type is used to indicate sub-types.

The Value Field consists of two sub-fields:

Global Administrator sub-field: 2 octets


This sub-field contains an Autonomous System number assigned by IANA.

Local Administrator sub-field: 4 octets

The organization identified by Autonomous System number in the Global Administrator sub-field can encode any information in this sub-field. The format and meaning of the value encoded in this sub-field should be defined by the sub-type of the community.

3.2. IPv4 Address Specific Extended Community

This is an extended type with Type Field composed of 2 octets and Value Field composed of 6 octets.



The value of the high-order octet of this extended type is either 0x01 or 0x41. The low-order octet of this extended type is used to indicate sub-types.

The Value field consists of two sub-fields:

Global Administrator sub-field: 4 octets


This sub-field contains an IPv4 unicast address assigned by one of the Internet registries.

Local Administrator sub-field: 2 octets

The organization that has been assigned the IPv4 address in the Global Administrator sub-field can encode any information in this sub-field. The format and meaning of this value encoded in this sub-field should be defined by the sub-type of the community.

3.3. Opaque Extended Community

This is an extended type with Type Field composed of 2 octets and Value Field composed of 6 octets.



The value of the high-order octet of this extended type is either 0x03 or 0x43. The low-order octet of this extended type is used to indicate sub-types.

This is a generic community of extended type. The value of the sub-type that should define the Value Field is to be assigned by IANA.

4. Route Target Community

The Route Target Community identifies one or more routers that may receive a set of routes (that carry this Community) carried by BGP. This is transitive across the Autonomous System boundary.

The Route Target Community is of an extended type.

The value of the high-order octet of the Type field for the Route Target Community can be 0x00, 0x01, or 0x02. The value of the low-order octet of the Type field for this community is 0x02.

When the value of the high-order octet of the Type field is 0x00 or 0x02, the Local Administrator sub-field contains a number from a numbering space that is administered by the organization to which the Autonomous System number carried in the Global Administrator sub-field has been assigned by an appropriate authority.

When the value of the high-order octet of the Type field is 0x01, the Local Administrator sub-field contains a number from a numbering space that is administered by the organization to which the IP address carried in the Global Administrator sub-field has been assigned by an appropriate authority.

One possible use of the Route Target Community is specified in [RFC4364].

5. Route Origin Community

The Route Origin Community identifies one or more routers that inject a set of routes (that carry this Community) into BGP. This is transitive across the Autonomous System boundary.

The Route Origin Community is of an extended type.

The value of the high-order octet of the Type field for the Route Origin Community can be 0x00, 0x01, or 0x02. The value of the low-order octet of the Type field for this community is 0x03.

When the value of the high-order octet of the Type field is 0x00 or 0x02, the Local Administrator sub-field contains a number from a numbering space that is administered by the organization to which the Autonomous System number carried in the Global Administrator sub-field has been assigned by an appropriate authority.

When the value of the high-order octet of the Type field is 0x01, the Local Administrator sub-field contains a number from a numbering space that is administered by the organization to which the IP address carried in the Global Administrator sub-field has been assigned by an appropriate authority.

One possible use of the Route Origin Community is specified in [RFC4364].

6. Operations

A BGP speaker may use the Extended Communities attribute to control which routing information it accepts or distributes to its peers.

The Extended Community attribute MUST NOT be used to modify the BGP best path selection algorithm in a way that leads to forwarding loops.

A BGP speaker receiving a route that doesn't have the Extended Communities attribute MAY append this attribute to the route when propagating it to its peers.

A BGP speaker receiving a route with the Extended Communities attribute MAY modify this attribute according to the local policy.

By default if a range of routes is to be aggregated and the resultant aggregates path attributes do not carry the ATOMIC_AGGREGATE attribute, then the resulting aggregate should have an Extended Communities path attribute that contains the set union of all the Extended Communities from all of the aggregated routes. The default behavior could be overridden via local configuration, in which case handling the Extended Communities attribute in the presence of route aggregation becomes a matter of the local policy of the BGP speaker that performs the aggregation.

If a route has a non-transitivity extended community, then before advertising the route across the Autonomous System boundary the community SHOULD be removed from the route. However, the community SHOULD NOT be removed when advertising the route across the BGP Confederation boundary.

A route may carry both the BGP Communities attribute, as defined in [RFC1997]), and the Extended BGP Communities attribute. In this case, the BGP Communities attribute is handled as specified in [RFC1997], and the Extended BGP Communities attribute is handled as specified in this document.

7. IANA Considerations

All the BGP Extended Communities contain a Type field. The IANA has created a registry entitled, "BGP Extended Communities Type". The IANA will maintain this registry.

The Type could be either regular or extended. For a regular Type the IANA allocates an 8-bit value; for an extended Type the IANA allocates a 16-bit value.

The value allocated for a regular Type MUST NOT be reused as the value of the high-order octet when allocating an extended Type. The value of the high-order octet allocated for an extended Type MUST NOT be reused when allocating a regular Type.

The Type field indicates where the Extended Community is transitive or not. Future requests for assignment of a Type value must specify whether the Type value is intended for a transitive or a non-transitive Extended Community.

Future assignment are to be made using either the Standards Action process defined in [RFC2434], the Early IANA Allocation process defined in [RFC4020], or the "First Come First Served" policy defined in [RFC2434].

The following table summarizes the ranges for the assignment of Types:

Type	Standard Action Early IANA Allocation	First Come First Served
-----	-----	-----
regular, transitive	0x90-0xbf	0x00-x3f
regular, non-transitive	0xd0-0xff	0x40-0x7f
extended, transitive	0x9000-0xbfff	0x0000-0x3fff
extended, non-transitive	0xd000-0xffff	0x4000-0x7fff

Assignments consist of a name and the value.

The Type values 0x80-0x8f and 0xc0-0xcf for regular Types, and 0x8000-0x8fff and 0xc000-0xcfff for extended Types are for Experimental use as defined in RFC 3692.

This document defines a class of extended communities called two-octet AS specific extended community for which the IANA is to create and maintain a registry entitled "Two-octet AS Specific Extended Community". All the communities in this class are of extended Types. Future assignment are to be made using the "First Come First Served" policy defined in [RFC2434]. The Type values for the transitive communities of the two-octet AS specific extended community class are 0x0000-0x00ff, and for the non-transitive communities of that class are 0x4000-0x40ff. Assignments consist of a name and the value.

This document makes the following assignments for the two-octet AS specific extended community:

Name	Type	Value
----	-----	-----
two-octet AS specific Route Target	0x0002	
two-octet AS specific Route Origin	0x0003	

This document defines a class of extended communities called IPv4 address specific extended community for which the IANA is to create and maintain a registry entitled "IPv4 Address Specific Extended Community". All the communities in this class are of extended Types. Future assignment are to be made using the "First Come First Served" policy defined in [RFC2434]. The Type values for the transitive communities of the two-octet AS specific extended community class are 0x0100-0x01ff, and for the non-transitive communities of that class are 0x4100-0x41ff. Assignments consist of a name and the value.

This document makes the following assignments for the IPv4 address specific extended community:

Name	Type	Value
----	-----	-----
IPv4 address specific Route Target	0x0102	
IPv4 address specific Route Origin	0x0103	

This document defines a class of extended communities called opaque extended community for which the IANA is to create and maintain a registry entitled "Opaque Extended Community". All the communities in this class are of extended Types. Future assignment are to be made using the "First Come First Served" policy defined in [RFC2434]. The Type values for the transitive communities of the opaque extended community class are 0x0300-0x03ff, and for the non-transitive communities of that class are 0x4300-0x43ff. Assignments consist of a name and the value.

When requesting an allocation from more than one registry defined above, one may ask for allocating the same Type value from these registries. If possible, the IANA should accommodate such requests.

8. Security Considerations

This extension to BGP has similar security implications as BGP Communities [RFC1997].

This extension to BGP does not change the underlying security issues. Specifically, an operator who is relying on the information carried in BGP must have a transitive trust relationship back to the source of the information. Specifying the mechanism(s) to provide such a relationship is beyond the scope of this document.

9. Acknowledgements

The authors would like to thank John Hawkinson, Jeffrey Haas, Bruno Rijsman, Bill Fenner, and Alex Zinin for their suggestions and feedback.

10. Normative References

- [BGP-4] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, August 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [RFC4020] Kompella, K. and A. Zinin, "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 4020, February 2005.

11. Informative References

- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

Authors' Addresses

Srihari R. Sangli
Cisco Systems, Inc.

E-Mail: rsrihari@cisco.com

Dan Tappan
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824

E-Mail: tappan@cisco.com

Yakov Rekhter
Juniper Networks, Inc.
1194 N. Mathilda Ave
Sunnyvale, CA 94089

E-Mail: yakov@juniper.net

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).