

Network Working Group
Internet Draft
Expiration Date: December 2001

Srihari Ramachandra
Daniel Tappan
Cisco Systems

Yakov Rekhter
Juniper Networks

BGP Extended Communities Attribute

[draft-ramachandra-bgp-ext-communities-09.txt](#)

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) except that the right to produce derivative works is not granted.

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>.

2. Abstract

This document describes an extension to BGP [[BGP-4](#)] which may be used to provide flexible control over the distribution of routing information.

3. Introduction

The Extended Community Attribute provides two important enhancements over the existing BGP Community Attribute:

- It provides 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 application 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 Autonomous system boundary. Without structure this can only be accomplished by explicitly enumerating all community values which will be denied or allowed and passed to BGP speakers in neighboring ASes based on the transitive property.

4. BGP Extended Communities Attribute

The Extended Communities Attribute is a transitive optional BGP attribute. The attribute consists of a set of "extended communities". Each extended community is coded as an eight octet value. All routes with the Extended Communities attribute belong to the communities listed in the attribute.

The Extended Communities Attribute has Type Code 16.

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

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

Type Field:

The value of the high-order octet will determine if its a regular type or 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 high-order octet of the Type Field is as shown below:

First bit (MSB)	:	IANA authority bit
		Value 0 : IANA assignable type
		Value 1 : Vendor-specific types
Second bit	:	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

For regular types, values of the Type Field 0 through 0x3f inclusive are assignable by IANA. The values 0x80 through 0xbf

inclusive are vendor-specific.

For extended types, values of the Type Field 0 through 0x3fff inclusive are assignable by IANA. The values 0x8000 through 0xbfff inclusive are vendor-specific.

Value Field:

The encoding of the Value Field is dependent on the "type" of the community as specified by the Type Field. The encoding of the community for the transitive communities should be such that it is unique globally (i.e. across the Autonomous Systems).

Two extended communities are declared equal only when entire 8 octets are equal.

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

5. New BGP Extended Community Types.

This document introduces few extended types and defines the Value Field for those types.

Type 0x00:

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

The value of the high-order octet of this extended type is 0x00. The low-order octet of this extended type is used to indicate subtypes.

The Value Field consists of two subfields:

Global Administrator subfield: 2 octets

This subfield contains an Autonomous System number assigned by IANA.

Local Administrator subfield: 4 octets

The organization identified by Autonomous System number in the Global Administrator subfield, can encode any information in this subfield. The value and meaning of the value encoded in this subfield should be defined by the subtype of the community.

Type 0x01:

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

The value of the high-order octet of this extended type is 0x01. The low-order octet of this extended type is used to indicate subtypes.

The Value field consists of two subfields.

Global Administrator subfield: 4 octets

This subfield contains an IPv4 address assigned by IANA.

Local Administrator subfield: 2 octets

The organization which has been assigned the IPv4 address in the Global Administrator subfield, can encode any information in this subfield. The value and meaning of this value encoded in this subfield should be defined by the subtype of the community.

Type 0x02:

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

The value of the high-order octet of this extended type is 0x02. The low-order octet of this extended type is used to indicate subtypes.

The Value Field consists of two subfields.

Global Administrator subfield: 4 octets

This subfield contains a 4-octets Autonomous System number assigned by IANA.

Local Administrator subfield: 2 octets

The organization identified by Autonomous System number in the Global Administrator subfield, can encode any information in this subfield. The value and meaning of the value encoded in this subfield should be defined by the subtype of the community.

6. 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 value of the Type field for the Route Target Community is 0x00 or 0x01. The value of the low-order octet of the extended type field for this community is 0x02.

When the value of the Type field is 0x00, the value of the Local Administrator subfield in the Value Field MUST be unique within the Autonomous system carried in the Global Administrator subfield.

7. 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 value of the Type field for the Route Origin Community is 0x00 or 0x01. The value of the low-order octet of the extended type field for this community is 0x03.

When the value of the Type field is 0x00, the value of the Local Administrator subfield in the Value Field MUST be unique within the Autonomous system carried in the Global Administrator subfield.

8. Link Bandwidth Community

When a router 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 non-transitive across the Autonomous system boundary.

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 MUST represent the Autonomous System of the router that attaches the Link Bandwidth Community. When a router receives a route with the community, the router may check the AS number in the Global Administrator subfield to see if its if its not the local AS and hence ignore the information carried in the Link Bandwidth Community.

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.

9. Operations

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

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.

A BGP speaker should not propagate the non-transitive extended community across the Autonomous system 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.

10. IANA Considerations

The Type Field values 0x00 and 0x01 are assigned in this document. Type Field values 2-0x3fff for extended-types (2-0x3f for regular types) are to be assigned by IANA, using the "First Come First Served" policy defined in [RFC 2434](#). Type values 0x8000-0xbfff for extended-types (0x80-0xbf for regular-types) are for vendor-specific types, and values in this range are not to be assigned by IANA.

11. Security Considerations

This extension to BGP does not change the underlying security issues.

12. Acknowledgements

The authors would like to thank John Hawkinson, Jeffrey Haas for their feedback.

13. References

[BGP-4] Rekhter, Y., and T. Li, "A Border Gateway Protocol 4 (BGP-4)", [RFC 1771](#), March 1995.

[RFC1997] Chandra, R., Traina, P., Li, T., "BGP Communities Attribute", [RFC1997](#), August 1996.

14. Author Information

Srihari Ramachandra
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
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