

Network Working Group
Internet Draft
Expiration Date: September 2004

Srihari R. Sangli (Procket Networks)
Daniel Tappan (Cisco Systems)
Yakov Rekhter (Juniper Networks)

BGP Extended Communities Attribute

[draft-ietf-idr-bgp-ext-communities-07.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](#).

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-4 which may be used to provide flexible control over the distribution of routing information.

3. 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](#)].

4. 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 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 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.

5. 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 eight octet quantity, as follows:

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

policy and experimental

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

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.

encoded in this sub-field should be defined by the sub-type of the community.

6.3. Four-octet AS specific extended community

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 either 0x02 or 0x42. 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 a 4-octets Autonomous System number assigned by IANA.

Local Administrator sub-field: 2 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.

6.4. Opaque extended community

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 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 which should define the Value Field is to be assigned by IANA.

7. 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 which is administered by the organization to which the Autonomous System number carried in the Global Administrator subfield 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 which is administered by the organization to which the IP address carried in the Global Administrator subfield has been assigned by an appropriate authority.

8. 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 which is administered by the organization to which the Autonomous System number carried in the Global Administrator subfield 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 which is administered by the organization to which the IP address carried in the Global Administrator subfield has been assigned by an appropriate authority.

9. 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 Link Bandwidth Community is of an extended type.

The value of the high-order octet of the Type Field is 0x00. The value of the low-order octet of the Type field for this community is 0x04.

The value of the Global Administrator sub-field in the Value Field MUST represent the Autonomous System of the router that attaches 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 sub-field of the Value Field.

10. 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 which 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.

11. IANA Considerations

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 value of the high-order octet of the Type Field defined as regular type (or extended type) for a extended community MUST NOT be reused as the value of the high-order octet of the Type Field defined as extended type (or regular type).

For the high-order octet of the Type Field, values 0x00 through 0x03 and 0x40 through 0x43 are assigned in this document and are defined as extended types.

For the combination of the high-order and low-order octets of the Type Field values 0x0002-0x0004, 0x0102-0x0103, and 0x0202-0x0203 are assigned in this document.

The Type Field values 0x04-0x3f and 0x44-0x7f for regular types (0x0400-0x3fff and 0x4400-0x7fff when expressed as extended types) are to be assigned by IANA, using the "First Come First Served" policy defined in [RFC 2434](#).

The extended Type Field values 0x0000-0x0001, 0x0005-0x00ff, 0x0100-0x0101, 0x0104-0x01ff, 0x0200-0x0201, 0x0204-0x02ff, 0x0300-0x03ff, and 0x4000-0x43ff are to be assigned by IANA, using the "First Come First Served" policy defined in [RFC 2434](#).

The Type Field values 0x90-0xbf and 0xd0-0xff for regular types (0x9000-0xbfff and 0xd000-0xffff when expressed as extended types) are to be assigned by IANA, using the "IETF Consensus" policy defined in [RFC2434](#).

The Type Field values 0x80-0x8f and 0xc0-0xcf for regular types (0x8000-0x8fff and 0xc000-0xcfff when expressed as extended types) are experimental, and are not to be assigned by IANA.

12. Security Considerations

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

13. Acknowledgements

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

14. Normative 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.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

15. Author Information

Srihari R. Sangli
Procket Networks, Inc.
1100 Cadillac Court
Milpitas, CA - 95035
e-mail: srihari@procket.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