

Network Working Group
Request for Comments: 1997
Category: Standards Track

R. Chandra
P. Traina
cisco Systems
T. Li
August 1996

BGP 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.

Abstract

Border Gateway Protocol [1] is an inter-autonomous system routing protocol designed for TCP/IP internets.

This document describes an extension to BGP which may be used to pass additional information to both neighboring and remote BGP peers.

The intention of the proposed technique is to aid in policy administration and reduce the management complexity of maintaining the Internet.

Introduction

BGP supports transit policies via controlled distribution of routing information. Mechanisms for this are described in [1] and have been successfully used by transit service providers. However, control over the distribution of routing information is presently based on either IP address prefixes or on the value of the AS_PATH attribute (or part of it).

To facilitate and simplify the control of routing information this document suggests a grouping of destinations so that the routing decision can also be based on the identity of a group. Such a scheme is expected to significantly simplify a BGP speaker's configuration that controls distribution of routing information.

Terms and Definitions

Community

A community is a group of destinations which share some common property.

Each autonomous system administrator may define which communities a destination belongs to. By default, all destinations belong to the general Internet community.

Examples

A property such as "NSFNET sponsored/AUP" could be added to all AUP compliant destinations advertised into the NSFNET. NSFNET operators could define a policy that would advertise all routes, tagged or not, to directly connected AUP compliant customers and only tagged routes to commercial or external sites. This would insure that at least one side of a given connection is AUP compliant as a way of enforcing NSF transit policy guidelines.

In this example, we have just eliminated the primary motivation for a complex policy routing database that is used to generate huge prefix and AS path based filter rules. We have also eliminated the delays caused by the out-of-band maintenance of this database (mailing in NACRs, weekly configuration runs, etc.)

A second example comes from experience with aggregation. It is often useful to advertise both an aggregate prefix and the component more-specific prefixes that were used to form the aggregate to optimize "next hop" routing. These component prefixes are only useful to the neighboring BGP peer or perhaps the autonomous system of the neighboring BGP peer, so it is desirable to filter this information. By specifying a community value that the neighboring peer or peers will match and filter on, these more specific routes may be advertised with the assurance that they will not propagate beyond their desired scope.

COMMUNITIES attribute

This document creates the COMMUNITIES path attribute is an optional transitive attribute of variable length. The attribute consists of a set of four octet values, each of which specify a community. All routes with this attribute belong to the communities listed in the attribute.

The COMMUNITIES attribute has Type Code 8.

Communities are treated as 32 bit values, however for administrative assignment, the following presumptions may be made:

The community attribute values ranging from 0x00000000 through 0x0000FFFF and 0xFFFF0000 through 0xFFFFFFFF are hereby reserved.

The rest of the community attribute values shall be encoded using an autonomous system number in the first two octets. The semantics of the final two octets may be defined by the autonomous system (e.g. AS 690 may define research, educational and commercial community values that may be used for policy routing as defined by the operators of that AS using community attribute values 0x02B20000 through 0x02B2FFFF).

Well-known Communities

The following communities have global significance and their operations shall be implemented in any community-attribute-aware BGP speaker.

NO_EXPORT (0xFFFFFFFF01)

All routes received carrying a communities attribute containing this value MUST NOT be advertised outside a BGP confederation boundary (a stand-alone autonomous system that is not part of a confederation should be considered a confederation itself).

NO_ADVERTISE (0xFFFFFFFF02)

All routes received carrying a communities attribute containing this value MUST NOT be advertised to other BGP peers.

NO_EXPORT_SUBCONFED (0xFFFFFFFF03)

All routes received carrying a communities attribute containing this value MUST NOT be advertised to external BGP peers (this includes peers in other members autonomous systems inside a BGP confederation).

Operation

A BGP speaker may use this attribute to control which routing information it accepts, prefers or distributes to other neighbors.

A BGP speaker receiving a route that does not have the COMMUNITIES path attribute may append this attribute to the route when propagating it to its peers.

A BGP speaker receiving a route with the COMMUNITIES path attribute may modify this attribute according to the local policy.

Aggregation

If a range of routes is to be aggregated and the resultant aggregates attribute section does not carry the ATOMIC_AGGREGATE attribute, then the resulting aggregate should have a COMMUNITIES path attribute which contains all communities from all of the aggregated routes.

Applicability

The COMMUNITIES path attribute may be used with BGP version 2 and all subsequent versions of BGP unless specifically noted otherwise.

Security Considerations

Security issues are not discussed in this memo.

Acknowledgments

We'd like to thank Vince Fuller, Sean Doran, and Andrew Partan for bringing to our attention the problems that we believe the BGP communities attribute will help solve. We'd also like to thank Yakov Rekhter his review of this document as well as his constructive and valuable comments.

Authors' Addresses

Paul Traina
cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134

E-Mail: pst@cisco.com

Ravishanker Chandrasekeran
(Ravi Chandra)
cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134

E-Mail: rchandra@cisco.com

Tony Li
E-Mail: tli@skat.usc.edu

References

- [1] RFC 1771
Rekhter, Y., and T. Li, "A Border Gateway Protocol 4 (BGP-4)",
March 1995.
- [2] RFC 1965
Traina, P., "Autonomous System Confederations for BGP", June 1996.