

Network Working Group
Request for Comments: 1998
Category: Informational

E. Chen
MCI
T. Bates
Cisco Systems
August 1996

An Application of the BGP Community Attribute in Multi-home Routing

Status of This Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document presents an application of the BGP community attribute [2] in simplifying the implementation and configuration of routing policies in the multi-provider Internet. It shows how the community based configuration can be used to replace the AS-based customization of the BGP "LOCAL_PREF" attribute, a common method used today. Not only does the technique presented simplify configuration and management at the provider level, it also represents a paradigm shift in that it gives the potential for the customer to control its own routing policy with respect to its service provider, as well as providing the ability for policy configuration to be done at a prefix based granularity rather than the more common AS based granularity.

1. Introduction

In the multi-provider Internet, it is common for a service subscriber (i.e., customer) to have more than one service provider, or to have arrangements for redundant connectivity to the global connected Internet. As discussed in [3], routing strategies in these cases usually require coordination between the service subscriber and its providers, which typically leads to customization of router configurations (e.g., BGP "LOCAL_PREF") not only by the subscriber, but also by its providers. Due to the large number of customers a provider serves, customization of router configurations at the provider level may present management and scalability problems.

This document presents an application of the BGP community attribute in simplifying the implementation of routing strategies in the multi-provider Internet. More specifically, the technique presented uses a community-based, rather than the common AS-based,

configuration of the BGP "LOCAL_PREF". It essentially removes the need for customized configuration of the BGP "LOCAL_PREF" attribute at the provider level while maintaining the same level of routing functionality and flexibility.

It also represents a paradigm shift in that it gives the potential for the customer to control its own routing policy with respect to its service provider, as well as providing the ability for policy configuration to be done at a prefix based granularity rather than the more common AS based granularity in use today.

2. AS-based Configuration and its Drawbacks

As discussed in [3], in today's multi-provider Internet, customized configuration of the BGP "LOCAL_PREF" attribute is often required to implement common routing strategies such as load-sharing or backup. There are two main reasons:

- o Lack of available implementations and deployment of routing software that supports the "Destination Preference Attribute" (DPA) as specified in [4].

DPA allows one to specify a globally transitive preference so that return traffic favors certain path. As discussed in [3], the attribute will be very useful in influencing route selection for routes with identical "LOCAL_PREF" and equal AS-path length.

- o In the multi-provider Internet, it is common for a provider to assign higher BGP "LOCAL_PREF" values for routes from its customers than from other service providers. This practice provides some degree of protection for its customer routes, and it facilitates implementation of certain routing strategies. It, however, also complicates other routing implementations such as backup arrangement, thus, requiring customized "LOCAL_PREF" configuration.

Figure 1 shows a typical case of a backup arrangement in the multi-provider Internet. In Figure 1, AS1 and AS2 are both providers, and AS3 and AS4 are customers of AS1 and AS2, respectively. AS3 has entered a bilateral agreement with AS4 to provide backup to each other. That is, AS3 would use its direct link to AS4 to reach only AS4 in the normal circumstance, and for transit in the case of a failure between AS3 and AS1. To realize this routing agreement, AS3 requests that its provider AS1 adjust its BGP "LOCAL_PREF" configuration so that AS1 reaches AS4 via AS2.




Figure 1: Typical Backup Scenario

Primarily due to scalability and management concerns, most providers only perform "LOCAL_PREF" customization based on ASs, not on IP prefixes. If IP prefix-based "LOCAL_PREF" configuration is needed, a technique known as the BGP AS-path manipulation can be used. However, it is currently only available in certain vendor's products.

There are several drawbacks with the the practice of AS-based BGP "LOCAL_PREF" configuration at the provider level:

- o The implementation tends to less efficient due to the process of coordination and configuration. More importantly, the process needs to be repeated each time a change (e.g., adding a new AS) occurs.
- o The AS-based customization complicates router configuration and increases complexity of network operation. It has become a serious scalability issue for providers.
- o It can not implement prefix-based configuration without the AS-path manipulation (i.e., using fake AS).
- o Keeping configuration up-to-date is some times problematic.

3. How the BGP Community Attribute Can Help

3.1 Overview of the Community Attribute

The BGP community path attribute is an optional transitive attribute of variable length [1,2]. The attribute consists of a set of four octet values, each of which specify a community. The community attribute values are encoded using an AS number in the first two octets, with the remaining two octets defined by the AS. As defined in [2], a community is a group of destinations (i.e. prefixes) that share some common attribute. Each destination can belong to multiple communities. All prefixes with the community attribute belong to the communities listed in the attribute.

The BGP community allows one to group a set of prefixes and perform routing decisions based on the identity of the group.

The well-known communities NO_EXPORT (0xFFFFFFFF01) and NO_ADVERTISE (0xFFFFFFFF02) are intuitive, and can be used for optimizing routing and for improving route aggregation.

3.2 Community-based Configuration

With the BGP community attribute [2], a provider can now use community-based, rather than AS-based, configuration of BGP "LOCAL_PREF". The provider first needs to coordinate with its customers a set of communities to be mapped to certain BGP "LOCAL_PREF" values. The provider can then apply a uniform BGP configuration to all its customers that would capture routes with the community values, and set up the appropriate BGP "LOCAL_PREF" values accordingly. A customer that requires customization in its provider BGP "LOCAL_PREF" configuration can simply send the appropriate community values in its routing announcements.

The major advantages of using this technique include:

- o The customer has full control in the process, which makes a lot of sense as the customer is in a position to have better understanding about its own topology and routing policy requirement.
- o The effect of route-based customization in BGP "LOCAL_PREF" configuration by providers can now be achieved, thus, removing the need of AS-Path manipulation in certain cases.
- o It addresses the scalability issue facing providers as it distributes the configuration work to the customer that requires customization.

4. A Real-World Implementation Example

MCI currently makes heavy use of the BGP "LOCAL_PREF" attribute value as part of its routing policy configuration process. Different BGP "LOCAL_PREF" values are assigned for routes from different sources. Table 1 details these values:

Category	LOCAL_PREF
Customer Routes	100
Customer backup Routes	90
Other ISP Routes	80
Customer-Provided backup	70

Table 1: Defined LOCAL_PREF Values

Note:

- o The value '100' is the default value used within our network configuration.
- o In most cases, the MED attribute set by a customer is sufficient for customer backup routes (e.g., T1 backs up T3). However, in certain cases configuration of "LOCAL_PREF" will still be necessary until the BGP DPA attribute is available.

To make use of the BGP community attribute, several community values (MCI's AS number: 3561 = 0x0DE9) have been defined that can be used by customers to tag routes so that the appropriate "LOCAL_PREF" values are configured. Table 2 lists the appropriate community attribute values (and the mappings of community to LOCAL_PREF):

community	LOCAL_PREF
3561:70 (0x0DE90046)	70
3561:80 (0x0DE90050)	80
3561:90 (0x0DE9005A)	90

Table 2: Community to LOCAL_PREF Mapping

A customer requiring MCI to configure BGP "LOCAL_PREF" values other than the default can tag their routes with the defined communities. The community values can be configured either based on an AS path list or an IP address access list. A cisco systems software specific configuration example is given in Appendix A to show how this can be achieved.

A uniform BGP configuration (see Appendix B, again cisco systems software specific) is applied by MCI to peers with customers that configure the appropriate "LOCAL_PREF" values based on the communities received.

This technique has been tested and is in use with several customers, and the response has been very positive. We are in the process of migrating all other customized BGP "LOCAL_PREF" configurations to this uniform community based configuration approach.

5. References

- [1] Rekhter, Y., and Li, T., "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.
- [2] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, August 1996.
- [3] Chen, E., and T. Bates, "Current Practice of Implementing Symmetric Routing and Load Sharing in the Multi-Provider Internet", Work in Progress.
- [4] Chen, E., and T. Bates, "Destination Preference Attribute for BGP", Work in Progress.
- [5] Chen, E., and T. Bates, "Application of the BGP Destination Preference Attribute in Implementing Symmetric Routing", Work in Progress.
- [6] cisco systems, cisco IOS Software Version 10.3 Router Products Configuration Guide (Addendum), May 1995.

6. Security Considerations

Security issues are not discussed in this memo.

7. Acknowledgments

The authors would specifically like to thank Ravi Chandra, Tony Li and Paul Traina of cisco systems for devising and implementing the community attribute.

8. Authors' Addresses

Enke Chen
MCI
2100 Reston Parkway
Reston, VA 22091

Phone: +1 703 715 7087
EMail: enke@mci.net

Tony Bates
cisco Systems
170 West Tasman Drive
San Jose, CA 95134

Phone: +1 408 527 2470
EMail: tbates@cisco.com

Appendix

These appendices list cisco systems software specific configuration examples for configuring communities, and for uniform route-map definition that sets up the appropriate "LOCAL_PREF" values based on the corresponding community values. These examples are given purely to show a working example of how the desired effect discussed in this document can be achieved. Please refer to [6] for more specific information on cisco configuration and syntax.

Appendix A. Community Configuration

The community values can be configured either based upon an AS path list or based an IP address access list. Here is an example that includes both cases:

```
!!
router bgp xxxx
neighbor x.x.x.x remote-as 3561
neighbor x.x.x.x filter-list 20 out
neighbor x.x.x.x route-map config-community out
neighbor x.x.x.x send-community
!
!!# match all
ip as-path access-list 1 permit .*
!
!!# list of customer ASs
ip as-path access-list 20 permit ^$
ip as-path access-list 20 permit ^64700_
ip as-path access-list 20 deny .*
!
!!# AS path based matching, backup for another ISPs customer
ip as-path access-list 40 permit _64710_
ip as-path access-list 40 permit _64711_
ip as-path access-list 40 deny .*
!
!!# route-map
route-map config-community permit 10
match as-path 40
set community 0x0DE90046
route-map config-community permit 20
match as-path 1
!
```

Note: The community can also be configured based on IP prefixes instead of AS numbers. For example,

```
!  
access-list 101 permit ip 192.160.154.0 0.0.0.0 255.255.255.0 0.0.0.0  
!  
route-map config-community permit 10  
match ip address 101  
set community 0x0DE90046  
route-map config-community permit 20  
match as-path 1  
!
```

Appendix B. Uniform Route-map Configuration

Here is the uniform route-map that can be used for all BGP customers:

```
!!# routes primary via another ISP  
ip community-list 70 permit 0x0DE90046  
ip community-list 70 deny  
!  
!!# routes also homed to another ISP, but with DPA or  
!!# AS-path length as the tie-breaker  
ip community-list 80 permit 0x0DE90050  
ip community-list 80 deny  
!  
!!# customer backup routes  
ip community-list 90 permit 0x0DE9005A  
ip community-list 90 deny  
!  
!!# the route-map applied to BGP customers  
route-map set-customer-local-pref permit 10  
match community 70  
set local-preference 70  
route-map set-customer-local-pref permit 20  
match community 80  
set local-preference 80  
route-map set-customer-local-pref permit 30  
match community 90  
set local-preference 90  
route-map set-customer-local-pref permit 40  
match as-path 1  
set local-preference 100  
!
```

