

Destination Preference Attribute for BGP
<[draft-ietf-idr-bgp-dpa-04.txt](#)>

Status of this Memo

This document is an Internet Draft. 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. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress".

Please check the I-D abstract listing contained in each Internet Draft directory to learn the current status of this or any other Internet Draft.

Abstract

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

This document describes a new BGP path attribute termed "Destination Preference Attribute" (DPA) which can be used by a single autonomous system (AS) to specify globally transitive metrics in its routing announcement via BGP. The metric can then be used by upstream BGP speakers to favor certain path for return traffic. The application of this attribute includes facilitating the implementation of symmetric routing and load sharing in the multi-provider Internet.

Introduction

In certain cases there is a need for an autonomous system (AS) to specify a globally transitive preference in its routing announcement

via BGP so that the upstream BGP speakers can use the preference to favor certain path for return traffic. For instance, as discussed in [3], currently it is difficult to implement symmetric routing and load sharing in the multi-provider Internet due to the lack of this preference in BGP.

In this paper, we propose a new BGP attribute termed "Destination Preference Attribute" (DPA) to address such a need. More specifically, the DPA is a globally transitive metric that can be used by an AS to specify preference in its routing announcement so that the return traffic favors certain path. As illustrated in [4] through several examples, this metric, combined with AS-based "LOCAL_PREF" offers much greater flexibility and manageability in implementing symmetric inter-domain routing and load sharing in the multi-provider Internet.

Destination Preference Attribute (DPA)

This document proposes the DPA path attribute, which is an optional transitive attribute of fixed length. The attribute is represented by a pair <AS#, DPA value>. The AS# is a two octet non-negative integer, which denotes the AS that specifies the preference. The DPA value is a four octet non-negative integer.

The DPA attribute has Type Code 11.

Route Selection Process

The DPA attributes shall be used as a route selection criteria, after the "LOCAL_PREF" attribute is evaluated, and before the evaluation of the AS path length and the multi-exit-discriminator (MED) attribute. However, if a route contains both MED and DPA attributes from the same neighboring AS, the MED values shall be favored over DPA values for route selection.

The higher the DPA attribute value, the more preferred the route.

A route with missing DPA attribute must be treated as having an DPA attribute with value zero.

Operation

The AS that sets this attribute must include its AS number in the attribute. A BGP speaker may use the "LOCAL_PREF" attribute to select a different path other than the one specified by the DPA attribute value. This does not preclude an AS from re-setting this attribute. However, coordination with the upstream and/or downstream neighbors is strongly recommended.

To make sure that the MED attribute and not the DPA attribute is used in the selection of routes from multiple peers of the same neighboring AS, the DPA value, if set, must be identical for all peers with the same neighboring AS. It is an operational matter to ensure the correct setting of the DPA value for multiple peers to the same neighboring AS.

Aggregation

If aggregation is done, the resultant aggregate shall be treated as a new NLRI. No DPA attribute shall be derived from more specific NLRIs which formed the aggregate. The resultant aggregate is free to have the DPA attribute set if so desired.

Remarks

It is noted that this new BGP attribute is simple and requires little change to the current practice and operation of BGP4. Nevertheless, the new attribute would offer the flexibility of shifting more influence on route selection to where the route originates, which has become increasingly meaningful as the Internet becomes more complex and dynamic. At the same time, the autonomy of an AS is preserved as the "LOCAL_PREF" feature remains unchanged. A typical application of this attribute is illustrated in [\[4\]](#) where the DPA attribute is used to simplify the implementation of symmetric inter-domain routing and load-sharing.

Applicability

The DPA path attribute may be used with BGP version 4 and all subsequent versions of BGP unless specifically noted otherwise.

Security Considerations

Security considerations are not discussed in this memo.

Acknowledgments

The authors would like to thank Yakov Rekhter of cisco for his insightful comments and suggestions. We also acknowledge Ramesh Govindan (ISI) and Ravi Chandra (cisco) for their helpful comments.

References

- [1] Rekhter, Y., and Li, T., "A Border Gateway Protocol 4 (BGP-4)", [RFC1771](#), March 1995.
- [2] Y. Rekhter, and P. Gross, "Application of the Border Gateway Protocol in the Internet", [RFC1772](#), March 1995.
- [3] Chen, E., and Bates, T., "Current Practice of Implementing Symmetric Routing and Load Sharing in the Multi-Provider Internet", INTERNET-DRAFT, <[draft-ietf-idr-symm-multi-prov-02.txt](#)>, January 1996.
- [4] Chen, E., and Bates, T., "Application of the BGP Destination Preference Attribute in Implementing Symmetric Routing", INTERNET-DRAFT, <[draft-ietf-idr-dpa-application-02.txt](#)>, January 1996.
- [5] Antonov, V., "BGP AS Path Metrics", INTERNET DRAFT, <[draft-ietf-idr-bgp-metrics-00.txt](#)>, March 1995.
- [6] Rekhter, Y., "Routing in a Multi-provider Internet", [RFC1787](#), April 1995.

Author's Addresses

Enke Chen
MCI
2100 Reston Parkway
Reston, VA 22091

phone: +1 703 715 7087
email: enke@mci.net

Tony Bates
MCI
2100 Reston Parkway
Reston, VA 22091

phone: +1 703 715 7521
email: Tony.Bates@mci.net

