INTERNET-DRAFT
<draft-ietf-idr-dpa-application-01.txt>
Expires in six months

Enke Chen
Tony Bates
MCI

June 1995

Application of the BGP Destination Preference Attribute in Implementing Symmetric Routing

<draft-ietf-idr-dpa-application-01.txt>

Status of this Memo

Abstract

This paper presents applications of the proposed Destination Preference Attribute (DPA) for BGP. It shows how the DPA attribute can aid in the implementation of symmetric inter-domain routing in the multi-provider Internet.

1. Introduction

The Destination Preference Attribute (DPA) is proposed in [4] for BGP. This attribute can be used by 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. This paper presents a typical application of this attribute. It illustrates how the DPA attribute facilitates the implementation of symmetric inter-domain routing and

load-sharing for the the typical cases presented in [3].

This paper assumes that in general an ISP treats other ISPs equally (in terms of the "local_pref" parameter) in the route selection process. It also assumes the following order of preference is followed for the purpose of route selection: first the "local_pref" parameter, followed by the DPA attribute, the shortest AS-path, the MED and the IGP metric.

2. Application of the DPA Attribute

In [3] we present several typical topologies of Internet connections, their inter-domain routing requirements, and the current practice to implement these routing policies. This section illustrates how the DPA attribute can be used to facilitate the implementation.

2.1 An Entity with a Single Direct Provider

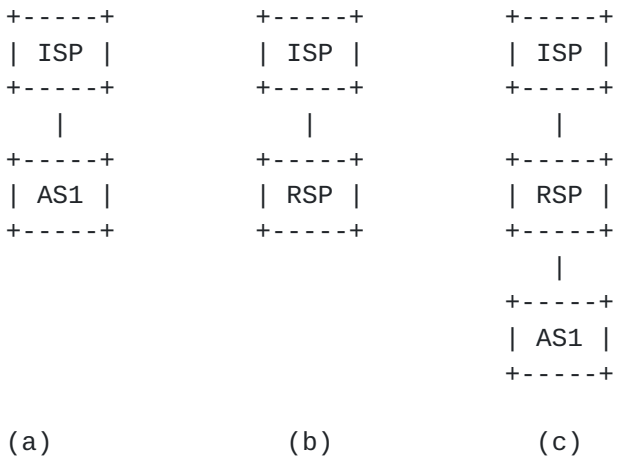

Figure 1

The routing is always symmetric at the inter-domain level. The DPA attribute should not be set as it is not needed.

AS1 can either take full routing or use default.

2.2 Backup of Entities with Different Direct Providers

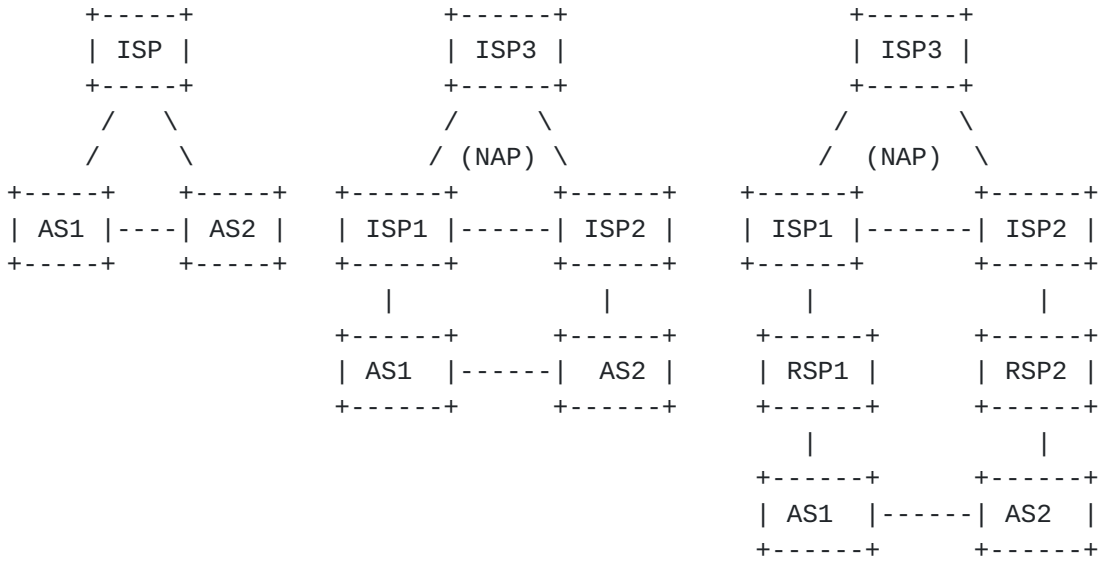

(a) (b) (c)

Figure 2

In all cases of Figure 2, in order to provide for backup, AS1 shall permit the acceptance of AS2's routes from both AS2 and AS1's direct providers, and permits their announcement to its direct providers. Similar configuration for AS2. As presented in [3], the AS-based "local_pref" configuration is sufficient to implement the routing requirements. It is not necessary to use the DPA attribute.

However, the DPA attribute would simplify the implementation as shown in the following.

Policy 1: Used solely as a backup link.

AS1 can simply announce all its routes with a higher DPA value to its direct provider, and with a lower DPA value to AS2. AS1 can either carry full routing or only take partial routing (AS2's routes) from both its direct provider and AS2, and configure default routes. Similar configuration for AS2.

Policy 2: Used for traffic between AS1 and AS2, and as backup in general

As with Policy 1, AS1 can simply announce all its routes with a higher DPA value to its direct provider, and with a lower DPA

value to AS2. AS1 also needs to configure proper AS-based "local_pref" value for AS2's routes. This is to override the higher DPA value for AS2's routes received from the direct provider. AS1 can either carry full routing or only take partial routing (AS2's routes) from both its direct provider and AS2, and configure default routes. Similar configuration for AS2.

2.3 An Entity with Multiple Direct Providers

This is where the DPA attribute would be most useful. As shown in Figure 3, AS1 has two direct providers. X and Y are routes of AS1. Note that AS1 could be an RSP.

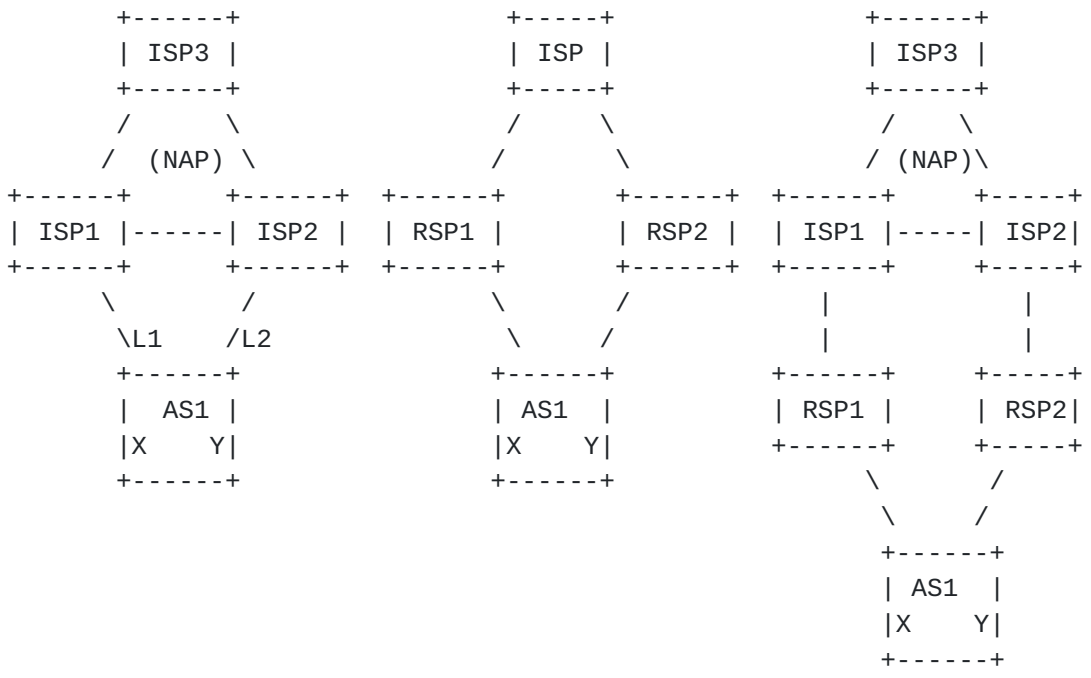

Figure 3

Policy 1: One link is used as primary, the other as pure backup

AS discussed in [3], the routing policy can be implemented by coordinating the AS-based "local_pref" parameter with direct providers. It is not necessary to use the DPA attribute.

However, the DPA attribute would simplify the implementation as detailed in the following.

AS1 can simply announce all its routes with a higher DPA value to the primary provider, and with a lower DPA value to the other provider.

AS1 can either carry full routing or use default. If AS1 takes full routing, then AS1 also need to configure AS-based "local_pref" so that the primary path is preferred.

Policy 2: Each link is used for traffic with the respective direct provider. In general one link is used as primary, the other as backup.

As with Policy 1, AS1 can simply announce all its routes with a higher DPA value to the primary provider, and with a lower DPA value to the other provider.

AS1 can be configured:

- o either with partial routing (only routes of the direct providers and their customers) and configure default routes with different weights.
- o or with full routing and configure AS-based "local_pref" values. The AS-list would still need to be updated and maintained, as discussed in [3].

Policy 3: Partial load-sharing among these links

That is, the direct link is used for traffic between AS1 and its direct providers including its customers. However, the closest exit point would be taken for traffic beyond these direct providers and their customers.

AS1 shall categorize its networks into two categories (say X and Y). Then, X routes shall be configured with higher DPA value when they are being announced to one direct provider, and with lower DPA values when they are being announced to the other direct provider. Similar configuration for Y routes.

In addition, AS1 also needs to configure AS-based "local_pref" so that the direct link is taken between the AS and its direct providers.

AS1 can take full routing. It can also take partial routing (routes of direct providers and their customers), and configure equal-weight default routes at its border routers and propagate them into its AS.

Policy 4: Complete load-sharing among these links

That is, each network in AS1 sends packets to the closer (in terms of internal route preference) border router that peers with a direct provider. The return traffic is expected to take a symmetric path.

AS1 shall categorize its networks into two categories (say X and Y). Then, X routes shall be configured with higher DPA value when they are being announced to one direct provider, and with lower DPA values when they are being announced to the other direct provider. Similar configuration for Y routes.

It is much simpler if AS1 does not take full routing. Then AS1 can configure default routes at its border routers and propagate them into its AS (via iBGP).

If AS1 still prefers to take full routing, this policy can only be achieved if AS1 manipulates the AS-path length so that routes received from the direct providers would have equal AS-path length. However, the routes with their AS paths manipulated would not be propagated upstream. That is, the propagation of the superfluous information in the AS path would be limited to the AS and possibly its downstream ASs, rather than the whole Internet.

3. Configure Preference for Routes with the DPA Attribute

It is possible, although not common, that the DPA attribute has been set by one AS (say AS1), and another AS (say AS2) desires further preference between its direct providers. The following options are available for AS2:

(1) AS2 uses the DPA attributes to do load sharing for routes other than AS1's. That is, AS2 does not include AS1's routes in load sharing with respect to AS2's direct providers. Instead, AS2 can coordinate with its direct providers to configure the proper AS-based "local_pref" values so that one provider is used as the primary, the other as the backup, for all of AS1's routes. This is to make sure that routing symmetry is maintained for routing to AS1. If there are multiple ASs that have configured the DPA attributes, then AS2 can perform load sharing by distribute (on per-AS basis) routes evenly with respect to its direct providers.

(2) AS1 chooses to re-set the DPA attribute for route announcements including AS1's routes. This may well cause the DPA attributes set by AS1 not to be used by upstream BGP speakers (due to non-comparable DPA attributes).

In many cases, Option (1) is probably preferred. However, Option (1) may not be able to maintain routing symmetry, either. It should be emphasized that when dealing with the complicated topologies of Internet connections, one needs to take into account its internal network topology, its connection to direct providers and to the major interconnection points. Coordination between providers is strongly recommended.

The following example illustrates Option (1). In Figure 4, AS1 has two direct providers RSP1 and RSP2. AS1 does load sharing by setting DPA attributes for routes W and Z. RSP1 has direct providers ISP1 and ISP2, and wishes to do load sharing.

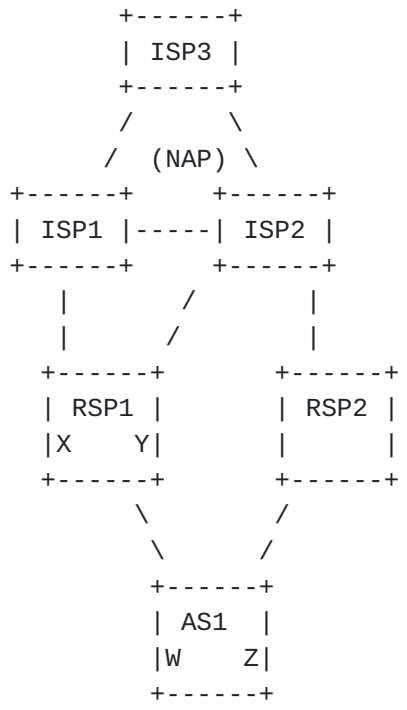

Figure 4

In this example, RSP1 can use the DPA attributes to do load sharing for routes without the DPA attributes. For AS1's routes (such as W and Z) that are already configured with the DPA attribute, RSP1 can coordinate, with ISP1 or ISP2, to configure the proper AS-based "local_pref" value so that one acts as primary to reach routes of AS1. For instance, ISP1 configures lower AS-based "local_pref" value for all of AS1's routes so that the ISP1 - ISP2 link is preferred to reach AS1's routes. This would also ensure that ISP3 would use ISP2 to reach AS1's routes.

4. Security Considerations

Security considerations are not discussed in this memo.

5. Acknowledgments

The authors would like to thank Yakov Rekhter of Cisco for his helpful comments and suggestions.

6. 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., "Analysis and Critique of the Current Practice of Implementing Symmetric Routing in the Multi-Provider Internet", INTERNET-DRAFT, <draft-ietf-idr-symm-multi-prov-01.txt>, June 1995.

[4] Chen, E., and Bates, T., "Destination Preference Attribute for BGP", INTERNET-DRAFT, <draft-ietf-idr-bgp-dpa-01.txt>, June 1995.

6. 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