Using a Dedicated AS for Sites Homed to a Single Provider <<u>draft-ietf-idr-as-dedicated-00.txt</u>>

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

Abstract

With the increased growth of the Internet, the number of customers using BGP4 has grown significantly. <u>RFC1930</u> outlines a set of guidelines for when one needs and should use an AS. However, the customer and service provider (ISP) are left with a problem as a result of this in that while there is no need for an allocated AS under the guidelines, certain conditions make the use of BGP4 a very pragmatic and perhaps only way to connect a customer homed to a single ISP. This paper proposes a solution to this problem in line with recommendations set forth in RFC1930.

1. Problems

With the increased growth of the Internet, the number of customers using BGP4 [1], [2] has grown significantly. <u>RFC1930</u> [4] outlines a set of guidelines for when one needs and should use an AS. However, the customer and service provider (ISP) are left with a problem as a result of this in that while there is no need for an allocated AS under the guidelines, certain conditions make the use of BGP4 a very pragmatic and perhaps only way to connect a customer homed to a single ISP. These conditions are as follows:

1) Customers multi-homed to single provider

Consider the scenario outlined in Figure 1 below.

	++	++
+	-+	
++	ISP A +	+ ISP B
Cust.++		
X +	-+	
++	++++\	++
		++
	++++ +-	-
	Cust.	ISP C
	Y	
	++	++

Figure 1: Customers multi-home to a single provider

Here both customer X and customer Y are multi-homed to a single provider, ISP A. Because these multiple connections are "localized" between the ISP A and its customers, the rest of the routing system (ISP B and ISP C in this case) doesn't need to see routing information for a single multi-homed customer any differently than a singly-homed customer as it has the same routing policy as ISP A relative to ISP B and ISP C. In other words, with respect to the rest of the Internet routing system the organization is singlyhomed, so the complexity of the multiple connections is not relevant in a global sense. Autonomous System Numbers (AS) are identifiers used in routing protocols and are needed by routing domains as part of the global routing system. However, as [4] correctly outlines, organizations with the same routing policy as their upstream provider do not need an AS.

Despite this fact, a problem exists in that many ISPs can only support the load-sharing and reliability requirements of a multihomed customer if that customer exchanges routing information

[Page 2]

using BGP-4 which does require an AS as part of the protocol.

2) Singly-homed customers requiring dynamic advertisement of NLRI's

While this is not a common case as static routing is generally used for this purpose, if a large amount of NLRI's need to be advertised from the customer to the ISP it is often administratively easier for these prefixes to be advertised using a dynamic routing protocol. Today, the only exterior gateway protocol (EGP) that is able to do this is BGP. This leads to the same problem outlined in condition 1 above.

As can be seen there is clearly a problem with the recommendations set forth in $[\underline{4}]$ and the practice of using BGP4 in the scenarios above. Section 2 proposes a solution to this problem with following sections describing the implications and application of the proposed solution.

It should also be noted that if a customer is multi-homed to more than one ISP then they are advised to obtain an official allocated AS from their allocation registry.

2. Solution

The solution we are proposing is that all BGP customers homed to the same single ISP use a single, dedicated AS specified by the ISP.

Logically, this solution results in an ISP having many peers with the same AS, although that AS exists in "islands" completely disconnected from one another.

Several practical implications of this solution are discussed in the next section.

3. Implications

3.1 Full Routing Table Announcement

The solution precludes the ability for a BGP customer using the dedicated AS to receive 100% full routes. Because of routing loop detection of AS path, a BGP speaker rejects routes with its own AS number in the AS path. Imagine Customer X and Customer Y maintain BGP peers with Provider A using AS number N. Then, Customer X will not be able to received routes of Customer Y. We do not believe that this would cause a problem for Customer X, though, because Customer X and Customer Y are both stub networks so default routing is adequate, and the absence of a very small portion of the full routing table is

[Page 3]

unlikely to have a noticeable impact on traffic patterns guided by MEDs received.

A BGP customer using the dedicated AS must carry a default route (preferably receiving from its provider via BGP).

3.2 Change of External Connectivity

The dedicated AS specified by a provider is purely for use in peering between its customers and the provider. When a customer using the dedicated AS changes its external connectivity, it may be necessary for the customer to reconfigure their network to use a different AS number (either a globally unique one if homed to multiple providers, or a dedicated AS of a different provider).

3.3 Aggregation

As BGP customers using this dedicated AS are only homed to one ISP, their routes allocated from its providers CIDR block do not need to be announced upstream by its provider as the providers will already be originating the larger block. [6].

<u>3.4</u> Routing Registries

The Internet Routing Registry (IRR) [5] is used by providers to generate route filtering lists. Such lists are derived primarily from the "origin" attribute of the route objects. The "origin" is the AS that originates the route. With multiple customers using the same AS, finer granularity will be necessary to generate the correct route filtering. For example, the "mntner" attribute or the "community" attribute of a route object can be used along with the "origin" attribute in generating the filtering lists.

4. Practice

The AS number specified by a provider can either be an AS from the private AS space (64512 - 65535) [4], or be an AS previously allocated to the provider. With the former, the dedicated AS like all other private AS's should be stripped from its AS path while the route is being propagated to the rest of the Internet routing system.

[Page 4]

<u>5</u>. Security Considerations

Security considerations are not discussed in this memo.

6. Acknowledgments

The authors would like to thank Roy Alcala of MCI and Arpakorn Boonkongchuen for their input to this document. The members of the IDR Working Group also provided helpful comments.

7. References

[1] Rekhter, Y., and Li, T., "A Border Gateway Protocol 4 (BGP-4)", <u>RFC1771</u>, March 1995.

[2] Rekhter, Y., and Gross, P., "Application of the Border Gateway Protocol in the Internet", <u>RFC1772</u>, March 1995.

[3] Rekhter, Y., "Routing in a Multi-provider Internet", <u>RFC1787</u>, April 1995.

[4] Hawkinson, J., and Bates, T., "Guidelines for creation, selection, and registration of an Autonomous System (AS)", <u>RFC1930</u>, March 1996.

 [5] T. Bates, E. Gerich, L. Joncheray, J-M. Jouanigot, D. Karrenberg,
M. Terpstra, & J. Yu., "Representation of IP Routing Policies in a Routing Registry (ripe-81++)", <u>RFC1786</u>, March 1995.

[6] E. Chen, J. Stewart., "A Framework for Inter-Domain Route Aggregation", <u>draft-ietf-idr-aggregation-framework-01.txt</u>, July 1997.

[Page 5]

8. Author's Addresses

John Stewart USC/ISI <u>4350</u> North Fairfax Drive Suite 620 Arlington, VA 22203 email: jstewart@isi.edu

Tony Bates Cisco Systems, Inc. <u>170</u> West Tasman Drive San Jose, CA 95134 email: tbates@cisco.com

Ravi Chandra Cisco Systems, Inc. <u>170</u> West Tasman Drive San Jose, CA 95134 email: rchandra@cisco.com

Enke Chen Cisco Systems, Inc. **170 West Tasman Drive** San Jose, CA 95134 email: enkechen@cisco.com

[Page 6]