

[draft-yu-simple-ipv6multihoming-00.txt](#)

A Simple Solution for IPv6 Multihoming

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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 and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

There is a challenge of providing IPv6 multihoming sites with the basic benefits such as redundancy and load sharing, and at the same time, scaling the global IPv6 routing table. It is equally challenging to provide a solution which is simple enough to be operational manageable. This document proposes a simple solution which meets the requirements and is also operationally manageable.


1. Purpose

There is a challenge of providing IPv6 multihoming sites with the basic benefits such as redundancy and load sharing, and at the same time, scaling the global IPv6 routing table. It is equally

challenging to provide a solution which is simple enough to be operational manageable. This document proposes a routing and addressing scheme that supports IPv6 multihoming which achieves the followings:

- a. Providing redundancy and load sharing for the multi-homed sites
- b. Scaling the global IPv6 Internet routing table
- c. Simple and operationally manageable

2. Proposal



A multi-homed customer will designate a primary ISP and receive IP address from its primary ISP's IPv6 aggregation block. In this example, ISP-1 is the primary ISP of customer-A and assigns Addr-1-A to the customer. Customer-A will advertise addr-1-A to ISP-1 and ISP-2 respectively. ISP-2 will advertise Addr-1-A to ISP-1 and to ISP-1 only. ISP-1 will, of course, advertise its own aggregation Addr-1 to the entire Internet.

For inbound traffic to Customer-A, traffic will first be forwarded to ISP-1 since it advertises Addr-1 which includes Addr-1-A. ISP-1 will then forward the traffic destined to Customer-A via its connection to ISP-2 and/or its direct link to customer-A, according to certain policies. Most common policy would be the use of the shortest exit. By using both connections to forward traffic to Customer-A, load sharing is achieved.

For outbound traffic from Customer-A, ISP-1 and ISP-2 can announce default and/or specific prefixes based on customer-A's requirements and traffic originated from Customer-A will be forwarded accordingly.

When the link between customer-A and ISP-2 (link-2 in the figure) goes down, all traffic will go in/out via the connection between Customer-A and ISP-1 (link-1 in the figure). Likewise, when link-1 is experiencing an outage, link-2 will be used for the traffic. This is because ISP-1 will continue announce its aggregate block (i.e. Addr-1) to the entire Internet and ISP-2 will still advertise Addr-1-A to ISP-1, all the inbound traffic to the customer will be take the path: ISP-1 -> ISP-2 -> Customer-A. Outbound traffic will automatically fall to link-2. This way, redundancy is achieved.

Same mechanism can be extended to those sites multihoming to more than two ISPs. Only those ISPs that the customer directly connected to would carry the more specific prefix assigned to the multi-homed customer.

The more specific prefix announcement could be covered under the peering agreement between ISPs as stated in [[RFC2546](#)].

3. Advantages of the proposal

- Achieves the goal of scaling the global routing table.
- It scales the global routing table. Only ISPs involved in the multihoming attachment need to know the more specific route of Addr-1. Other ISPs such as ISP-3 and ISP-4 in figure 1 do not need to know specifics.

If the two involved ISPs has no direct connection, the more specific route would need to be carried by other ISPs in the path. This seems to be lesser a problem since ISP assigned with an Internet visible aggregate block usually has direct connection to each other.

- It is really simple and thus more manageable and scalable operation wise. It does not require multiple addresses assigned all individual hosts of a multi-homed customer site, neither it requires $O(N^2)$ tunnels among different ISP routers.

4. Concerns

- The primary ISP of a multihoming customer (such as ISP-1 in the example) would need to do more work in terms of distributing the traffic among other connected ISPs and the customer link. This can be done using BGP policy (BGP community) and using shortest exit. It will also carry more traffic for the customer. However, this is a value added service to the customer and the primary ISP can charge more fees for such services.
- The primary ISP is the sole interface for the multihomed customer to the Internet other than the ISPs the customer has directly connection with. Outages such as one between ISP-1 and ISP-4 would impact the reachability from customers of ISP-4 to Customer-A even though ISP-2 still has good connection to ISP-4. However, if the primary ISP is a good quality ISP, this sort of situation should happen rarely. The reason is that it's common practice that an ISP, especially good quality ISP, has multiple connections to other big ISPs at different geographical locations. Good quality ISPs also have robust network design to prevent any

failure to impact the whole network. Choosing a good quality and robust ISP as primary ISP is a good practice.

It is desirable to have solution which provides perfect redundancy and, at the same time, simplicity to scale management and operation. In the absence of a perfect solution, the trade-off needs to be made. The author believes that it is very important that the solution has to be simple and manageable. Manageability should be the top consideration.

5. Security Considerations

Security considerations are out of scope of this document.

6. References

[RFC2546] A. Durand and B. Buclin, "6Bone Routing Practice."
[RFC2546](#), March 1999.<ftp://ftp.isi.edu/in-notes/rfc2546.txt>.

Author's Address

Jieyun (Jessica) Yu
UUNET Technology
880 Technology Dr.
Ann Arbor, MI 48108

Phone: (734) 214-7314

EMail: jyy@uu.net

----- End of Forwarded Message