IPv6 Multi-Homing Architectural Roadmap – Requirements Extract
Version 0.1

{Curly braces surround notes or editorial comments as seen here. Normal font is used for reminders to the author and bold face font, as in this text, is intended for the reader, in particular, the mulsixhomers group.}

Status of this Memo

This document is a proposed working document of the ipv6mh mailing list. The ipv6mh web page is http://arneill-py.sacramento.co.us/ipv6mh.

Abstract

This document describes a framework and a transition path for IPv6 multi-homing that is consistent with aggregation goals of IPv6.

1.0 Introduction

IPv6 provides a significant challenge for sites using IPv4 multi-homing techniques. These techniques are generally used by sites seeking to

- provide survivability from failures of site exit routers, transit provider connectivity and transit providers themselves or
- load balance between two or more transit providers.

Failure survival can be of either:

- the Site's Internet presence (access to and by the Internet following the failure, not user transparent) or
- active Transport layer sessions (user transparent).

IPv4 address availability will influence the ability of an IPv4 site to achieve both types of survival. Sites with Provider Independent (PI) IPv4 addresses may achieve both types of survival. Site required to use NAT with each transit provider will be able to achieve the first, but not the second and load balancing is impossible. For sites with IPv4 PI addresses, both multi-homing objectives (load balancing and failure survival at both levels) were achieved in IPv4 by the unrestricted flow of a site's IPv4 PI address (the site's assigned public addresses) routing information, and thereby a site's IPv4 packets, to and from the site's transit providers. "IPv4-style" multi-homing therefore represents different capabilities for different sites. To the extent that a site's Internet presence surviving failures has value, a partial IPv6 multi-homing solution offering this level of capability has value.

This document describes both a multi-homing architecture for IPv6 and transitioning mechanisms to move from IPv4 styled multi-homing to a complete support of the architecture. Such transitioning mechanisms for multi-homing within IPv6 are just as necessary as are transition mechanisms from IPv4 to IPv6 itself. That a site's Internet presence surviving failures has value, a partial IPv6 multi-homing solution offering this level of capability also has value. Just as a site's IPv4 address availability influences what it can achieve with IPv4 multi-homing, the address availability will also be a driver for site's moving to IPv6. Hence an IPv6 multi-homing solution that provides the same capability as the NAT IPv4 solution is desirable for the early adopters of IPv6. At no time is it required that the IPv6 Global Routing Table contain individual site prefixes. **{Better wording needed because the 6to4 mechanism for sites with IPv4 PI addresses violates this principal.}** This is considered to be the fundamental flaw of IPv4 multi-homing and any return to it is considered unacceptable in this roadmap.

The roadmap consists of the following sections:

- An introduction (Section 1.0)
- Multi-homing Requirements (Section 2.0)
- Roadmap Features (Section 3.0)
- Feature Requirement Scorecard (Section 4.0)
- Considerations (Section 5.0) and
- Appendices containing background information.

1.1 Terminology

Basic IPv6	IPv6 host capabilities as defined by specifications current in November, 2001.
Home	The process of accommodating multiple PA addresses and ingress filtering by transit providers during the selection of source and destination IP addresses.
MA	Multi Address
MH IPv6	Hosts capable of supporting multi-homing directly as defined in Section 3.2.1.
Multi-homed Host	A host supporting more than one data link layer interface.
Multi-homed Site	A site connected to more than one transit provider.
Re-home	The process of accommodating a new selection of source and destination IP addresses where multiple PA addresses and ingress filtering by transit providers exist.
SA	Single Address
Single-homed Host	A host supporting one data link layer interface
Single-homed Site	A site connected to one transit provider

2.0 Multi-homing Requirements

Requirements for multi-homing derive from three sources:

- The multi6 IETF working group (Section 2.1)
- Other IETF policies (Section 2.2)
- And the framework (Section 2.3)

The additional requirements from the framework are intended to clarify direction and scope of the framework.

2.1 IETF Site Multi-Homing Architecture Requirements

In May 2002, the IETF Network Working Group issued an Internet-Draft, "Requirements for IPv6 Site-Multihoming Architectures," draft-ietf-multi6-multihoming-requirements-03 [10] (Multi6 Draft). The actual status of this Internet-Draft seems to be unclear. It is still available on the IETF Internet-Draft site even though it expired on October 30, 2002. This Internet-Draft, in addition to placing some constraints and requirements on an actual IPv6 multi-homing architecture, provides a definition of IPv4 multi-homing. While placing constraints on an actual IPv6 site multi-homing architecture, it does not attempt to translate the IPv4 capabilities described into a set of requirements within the IPv6 context itself. This document provides this IPv6 context in Section 1.3 and the IPv6 factors which imply them in Sections 1.2.1 and 1.2.2. In general the architecture described here adheres to the requirements of the Multi6 Draft. Differences will be identified.

2.1.2 Basic IPv6

The Multi6 Draft defines Basic IPv6 hosts as hosts "implementing RFC 2373, RFC 2460, RFC 2553 and other basic IPv6 specifications current in November 2001." While not necessarily benefiting from multi-homing, such hosts must continue to function as if they were in a single homed site. This definition of Basic IPv6 does not define which specifications are "basic" and which are not, but the identified date essentially could include all specifications current at the time of the writing of the Multi6 Draft.

An alternative definition of Basic IPv6 can be found in "IPv6 Node Requirements" [12].

2.1.3 Transport Protocol Requirements

Transport Protocols MUST be able to

- home during session state establishment and
- re-home during the life of the session without resetting session state or awareness of the re-homing operation on the part of the Application Layer.

This requirement is the "Homing Transparency" requirement. Further, Homing Transparency MUST be supported for hosts in session (client or server) having the following capability levels:

- Basic IPv6 and Basic IPv6
- Basic IPv6 and MH IPv6 and
- MH IPv6 and MH IPv6.

What constitutes session state is Transport Layer protocol specific. Homing Transparency for the following Transport Layer protocols, their respective session states and architectural support levels are defined in the following table.

Protocol	Session State	Support
TCP	Window and acknowledgement	MUST
SCTP	{To be determined}	MUST
UDP	none	MUST
IPSEC	Security associations	SHOULD
ICMP	none	{TBD}
Raw IP	Application dependent	MAY

{What about Mobility requirements or are they doing their own thing?}

2.1.4 Site Connectivity Requirements

In addition to the host capability levels requiring support for Homing Transparency, sites with different Internet connectivity must be supported as well, specifically, either or both of the sites involved in a Transport Session must be connected to the Internet via two or more transit providers. Combining this requirement with the host capability requirement creates the following five configurations which must support Homing Transparency:

1. Basic IPv6 (single homed site) and Basic IPv6 (multi-homed site)
2. Basic IPv6 (multi-homed site) and Basic IPv6 (multi-homed site)
3. Basic IPv6 (single-homed site) and MH IPv6 (multi-homed site)
4. Basic IPv6 (multi-homed site) and MH IPv6 (multi-homed site)
5. MH IPv6 (single or multi-homed site) and MH IPv6 (single or multi-homed site).

The configuration wherein both sites are single homed and supporting Basic IPv6 hosts is the current IPv6 environment. Homing Transparency is not at issue because it does not exist (hence the driving force for this framework). Transport sessions in this environment do not need to change source or destination addresses or succumb to ingress filtering issues.

Additionally, hosts supporting the framework, MH IPv6 hosts, can function in either single or multi-homed settings. A single homed site is simply a multi-homed site which has lost all but one of its transit providers.

2.2 IPv6 Policy Driven Requirements

The policies identified in Appendix A dictate certain requirements for IPv6 multi-homing solutions to provide functionality comparable to that of IPv4:

- Transport Layer session initiation MUST accommodate ingress filtering by transit providers through an awareness of routing paths and
- Transport Layer connections MUST be able to utilize different source and destination addresses during the life of the connection as routing paths change during re-homing events.

The first facilitates both ingress filtering and load balancing during normal operations. The second allows Transport Layer connections to transparently survive re-homing events. These two functions can be translated into two specific requirements:

- Source and destination routing path awareness during the life of a Transport Layer session (Routing Awareness) and
- Re-homing event recognition and handling based upon Routing Awareness (Re-homing Events).

Load balancing is assumed to be facilitated by routing table state. The Routing Awareness requirement implies load balancing under normal circumstances.

Christian Huitema and Iljitsch Beijnum have suggested the following sequence to achieve IPv6 multi-homing:

1. Start with multi-addressing, because it does not require any global infrastructure. Hey, with two IPv4 addresses a site can get two 6to4 prefixes right now.
2. Cross that. there are things that don't work with multi-addressing, notably address selection and ingress filtering. So, start with multi-homing, and develop host based address selection algorithms.
3. Study and solve the ingress filtering problem and develop standard ways for exit routers to inform hosts about the state of the exit.
4. We have multiple addresses, but we only select one at the beginning of the connection. Use the MIPv6 constructs or a derivative to make sure that connections survive the failure of any specific site exit.
5. Now, we have several addresses and we know how to juggle them. Add a virtual address, in some form of PI scheme. Use the MIPv6 construct to move connections from the potentially slower virtual scheme to a "PA" scheme.

Steps 1-3 achieve a site multi-homing Internet presence survival capability and steps 4 and 5 extend the capability to transport layer session survivability.

2.3 Framework Requirements

This section provides additional detail for the architecture.

2.3.1 Intranet Isolation

Under IPv4 PI multi-homing, Internet topology changes had no effect upon internal site communications. Under this architecture, Internet related failures must be isolated to systems participating in Internet communication. Specifically, re-homing MUST only be necessary for communication which extends beyond the site.

2.3.2 Application Layer Requirements

Applications utilizing an API compatible with RFC 2553 MUST NOT require changes when a host implements this architecture. Applications utilizing other transport layer functionality MAY require changes.

Basic IPv6 compatibility requires that applications utilizing RFC 2553 on Basic IPv6 hosts must function with hosts upgraded to this architecture. This requirement extends the compatibility to hosts which incorporate this architecture.

Appendix A IPv6 Policies

Various IPv6 policies complicate and restrict a site’s ability to directly apply these IPv4 techniques within IPv6 to achieve these same objectives. These policies are described in appendices A.1-A.3.

A.1 No IPv6 Provider Independent Addresses

> An end site is defined as an end user (subscriber) who has a business relationship with a service provider that involves:
>
> - that service provider assigning address space to the end user
> - that service provider providing transit service for the end user to other sites
> - that service provider carrying the end user's traffic.
> - that service provider advertising an aggregate prefix route that contains the end user's assignment.
>
> [IPv6 Address Allocation and Assignment Assignment (sic) Policy, June, 26 2002, APNIC, ARIN, RIPE – NCC]

Sites may only acquire IPv6 addresses from their ISP's who in turn receive them from a Regional or National Internet Registry. For traffic to be routed to a site by an ISP the site must utilize addresses assigned by the ISP. The multi-homing site then ends up with multiple IPv6 addresses assigned to it which become deployed upon any host requiring communication with the IPv6 Internet. While this policy has the intent of restricting the size of the Internet routing table, it increases the routing tables for the multi-homing site. Further the administrative burden related to these addresses in the form of router access-lists or firewall rules is correspondingly increased for a multi-homing site.

A.2 No Mixing of Transit Provider Assigned Addresses

> Leaf sites or pNLAs MUST only advertise to an upstream provider the prefixes assigned by that provider. Advertising a prefix assigned by another provider to a provider is not acceptable, and breaks the aggregation model. A site MUST NOT advertise a prefix from another provider to a provider as a way around the multi-homing problem. [RFC 2772, Section 4]

This breaks the ability to non-disruptively recover from router, connectivity or ISP failures. Under IPv4 the PI address would continue to exist in the Internet and routing convergence would allow transport layer connections to continue. Under this policy when any of these connectivity elements fail, convergence does not occur and the transport connections using the addresses assigned by the ISP associated with the failure will fail. The policy explicitly mentions that it may not be violated for the purpose of resolving this situation.

A.3 Ingress Filtering

> By restricting transit traffic which originates from a downstream network to known, and intentionally advertised, prefix(es), the problem of source address spoofing can be virtually eliminated… . [RFC 2267, Section 3] All providers of Internet connectivity are urged to implement filtering described in this document to prohibit attackers from using forged source addresses which do not reside within a range of legitimately advertised prefixes. [RFC 2267, Section 1]

While RFC 2267 is not applicable only to IPv6, it has no impact on IPv4 multi-homing techniques because IPv4 PI addresses are "known and intentionally advertised" to each ISP and therefore are not filtered. Under IPv6, this means that any transport connection, assuming the routing situation created by 1.2 is fixed, may still fail, because the wrong source addresses are being sent to an ISP.

This policy extends the multi-homing difficulties beyond transport connection survivability during failures to the ability to create transport connections involving multi-homed sites. Somehow a source and destination address must be chosen that will guarantee that the communication leaving a multi-homed site and that returning from a multi-homed site will not be blocked by either site's transit providers due to ingress filtering.

Although a solution which manages to change source and destination addresses consistent with the new routing path to accommodate the requirements of ingress filtering and failure recovery would allow normal transport connections to survive, it would be inadequate in the presence of a host-to-host IPSEC communication. Security associations are tied to the IP addresses used between the end points. To allow the transport connection to survive, new IPSEC security associations would need to be established on both hosts.

Document and Author

The latest version of this document is always available at:

`http://arneill-py.sacramento.ca.us/ipv6mh/`

The author can be reached at:

Harold Grovesteen

2003 Winter Sunday Way

Arlington, TX 76012-4940

Email: `h.grovsteen@attbi.com`

References

[1] "Geographically Aggregatable Provider Independent Address Space to Support Multihoming in IPv6", work in progress

[2] "Provider Internal Aggregation based on Geography to Support Multihoming in IPv6", work in progress

[3] "An IPv6 Provider-Independent Global Unicast Address Format", work in progress

[4] "Application and Use of the IPv6 Provider Independent Global Unicast Address Format", work in progress

[5] J. Hagino and H. Snyder, "IPv6 Multihoming Support at Site Exit Routers", RFC 3178, October 2001

[6] "Multi Homing Aliasing Protocol", work in progress

[7] "Extension Header for Site-Multi-homing support", work in progress

[8] "Host-Centric IPv6 Multihoming", work in progress

[9] "Host Identity Protocol", work in progress

[10] "Requirements for IPv6 Site-Multihoming Architectures", work in progress

[11] {Add RFC 3280 reference info here.}

[12] "IPv6 Node Requirements", work in progress

[13] R. Gilligan, S. Thomson, J. Bound, W. Stevens, "Basic Socket Interface Extensions for IPv6", RFC 2553, March 1999

[14] Application of the MIPv6 protocol to the multi-homing problem, work in progress